SYSTEM AND METHOD FOR POLICY-GATED AUTHENTICATION AND CRYPTOGRAPHIC HASHING

TECHNICAL FIELD

[0001] The present invention generally relates to digital media security and video infrastructure. Specifically, the present invention relates to a system and method for authenticating digital video content at the binary sample level using cryptographic identity records generated during the virtualisation and indexing process. Further, the invention introduces a universal adapter architecture capable of unifying disparate authentication signals into a single, policy-gated, immutable cryptographic record.

BACKGROUND OF THE INVENTION

[0002] The proliferation of deepfakes has transitioned from a theoretical concern into a quantifiable global crisis that threatens financial, political, and social stability. The scale of this issue is severe enough to demand attention from the video industry.

[0003] The financial damage caused by deepfake technology is immense and has been verified through significant real-world incidents.

[0004] Industry surveys indicate that business vulnerability is nearly universal, with 92% of businesses reporting some degree of financial loss attributable to deepfakes.

[0005] Deepfake technology acts as a force multiplier for traditional cybercrime, significantly enhancing fraudulent schemes such as Business Email Compromise (BEC). In 2023, deepfake-

enhanced BEC schemes were responsible for over USD 2.9 billion in losses. This demonstrates the capacity of synthetic media to increase the effectiveness and scale of financial fraud.

[0006] The insurance industry has identified a sharp escalation in synthetic media fraud, particularly regarding voice attacks. Insurers reported a staggering 475% increase in synthetic voice attacks throughout 2024. This trend has led to warnings of a deepfake coverage gap as standard cyber policies may not provide adequate protection against these increasingly sophisticated impersonation tactics.

[0007] Beyond financial impacts, deepfakes inflict damage on the heart of democratic processes and global social trust.

[0008] Current detection systems including Reality Defender or Sensity AI, utilize neural network classifiers to analyze decoded pixel and audio output for synthetic artifacts.

[0009] These methods are fundamentally reactive and retrospective, often losing the "arms race" against generative models as AI improves.

[0010] Furthermore, detection technologies operate at the perceptual level and lack the architectural capacity to function at the binary sample level or compose results with disparate authentication sources.

[0011] Provenance and Metadata Solutions like C2PA and Content Credentials embed tamper-evident metadata and manifest chains at the point of capture to track content history.

[0012] However, this metadata exists alongside the file packaging and can be easily stripped or lost during re-encoding or distribution through platforms that remove EXIF data. These systems protect distribution channels rather than the underlying content and lack the ability to incorporate secondary signals, such as PRNU scores or forensic detector results, into a unified sample-level binding.

[0013] Invisible watermarking techniques, such as Google's SynthID, embed digital signals directly into the pixel data of a file. While the signal lives inside the content, it is notoriously fragile and susceptible to destruction through standard media transformations like compression, cropping, or color grading. Because watermarking is applied at the monolithic file level, it cannot address the internal structure of the video or unify multiple authentication inputs.

[0014] Some manufactures utilize hardware-level cryptographic signatures to verify content at the moment of capture. While providing a strong hardware-based signal, these remain single-source solutions that cannot be composed with broader policy logic. These methods lack the sample-level hashing and real-time assembly-time verification required for granular structural integrity throughout the media lifecycle.

[0015] Existing methodologies fundamentally treat video as a sealed, finished file, attempting to attach verification from the outside after the content has been created.

[0016] None of these approaches provide a unified architecture capable of orchestrating multiple authentication methods under a configurable trust policy or binding the combined results to an immutable, binary sample-level record. The lack of a foundational layer to lock verification results into the fundamental structure of the video asset remains a critical vulnerability in the current media ecosystem

[0017] The present invention addresses the aforementioned drawbacks through a unique combination of multi-source policy-gated authentication at ingestion and cryptographic hashing at the binary sample level. By utilizing immutable registration and real-time verification during virtual video assembly, the system provides a structural solution that survives re-encoding and resists the generative AI arms race.

[0018] Any discussion of the prior art throughout the specification should in no way be considered as an admission that such prior art is widely known or forms part of common general knowledge in the field.

SUMMARY OF THE INVENTION

PROBLEMS TO BE SOLVED

[0019] It may be an advantage to develop a system and method that establishes a cryptographic record proving content is authentic and human-created media at the binary sample level.

[0020] It may be an advantage to develop an authentication architecture not dependent on a reactive arms race of detection technologies which are retrospective, probabilistic, and easily avoided by subsequent generations of AI generators.

[0021] It may be an advantage to develop an efficient solution not limited to external metadata that can be stripped, broken during re-encoding, or lost when uploaded to platforms that remove EXIF data.

[0022] It may be an advantage to develop a unified architecture that acts as a universal adapter, orchestrating multiple disparate authentication methods including C2PA, PRNU sensor fingerprinting, and AI detector ensembles within a single layer.

[0023] It may be an advantage to develop a system wherein the authentication policies are expressed as signed data objects, utilizing composable Boolean logic such as ALL OF, ANY OF, and N OF M.

[0024] It may be an advantage to develop a method for interrogating rendered media to extract complete structural information from internal tables, such as sample to chunk mapping and chunk offset records.

[0025] It may be an advantage to develop a system that computes cryptographic hashes of raw binary data within the MDAT section of a container.

[0026] It may be an advantage to develop a compound record, such as a Sample Identity Manifest (SIM), that anchors both individual sample hashes and an Authentication Provenance Bundle to an immutable registry.

[0027] It may be an advantage to develop real-time assembly-time hash verification that can block playback, signal unverified samples, or log results without interrupting the media stream.

[0028] It may be an advantage to develop a computationally efficient system that includes two components, a configurable Authentication Policy Engine that evaluates inbound footage against a declared trust policy using pluggable Authentication Source Adapters and a Sample Identity Manifest (SIM) that records cryptographic hashes of individual binary samples.

[0029] It may be an advantage to develop an authentication method where the record of authenticity remains persistent and accessible independently of the media file, ensuring that while the registry record survives re-encoding, any alteration to the binary samples during re-encoding is immediately detectable as a hash mismatch.

MEANS FOR SOLVING THE PROBLEM

[0030] The present invention is directed to a system and method for authenticating digital video content as human-created at the individual binary sample level, providing a foundational infrastructure layer for deepfake detection and content provenance.

[0031] The invention establishes the right to identify references, ensuring that media content is cryptographically validated through a policy-gated architecture before it is registered in a trust system and subsequently verified during assembly-time playback.

[0032] In an aspect of the present invention, a system for policy-governed content verification with pluggable authentication is provided, comprising an authentication policy engine to evaluate inbound content against Boolean requirements, pluggable authentication source adapter modules to return standardized results, and a virtualization engine incorporating a Sample Indexing Module to separate structural container data from media binary sample data.

[0033] The system further consists of a sample identity manifest registry to store signed identity records, a token validation module to govern assembly conditions, and an assembly-time verification module to compare real-time hashes against the registered manifest.

[0034] In another aspect of the present invention, a computer-implemented method for policy-gated sample-level content authentication comprises receiving a digital video file with a reference to a signed authentication policy, invoking pluggable adapters and evaluating the Boolean expression of the policy against the collected adapter results. Upon successful evaluation, the method involves interrogating the video file to extract structural metadata, indexing individual binary samples by byte offset and temporal position, selecting a subset of samples for hashing, generating a compound sample identity manifest and authentication provenance bundle and registering these in an immutable registry to enable subsequent real-time verification during playback.

[0035] In a further aspect of the present invention, an integrated content governance and authenticated assembly framework is established where the virtualization of video into addressable building blocks allows for the binding of authentication provenance to the fundamental structural data of the media. This framework integrates with reference-based content architectures and distributed

ledger governance to create a complete stack that virtualizes structure, records access, governs assembly, and authenticates identity at the sample level.

[0036] Preferably, the system utilizes an Authentication Policy Engine that loads policies as signed data objects, expressing authentication requirements as Boolean logic over a plurality of authentication source types.

[0037] Preferably, the system satisfies the trust requirements through composable operators including ALL OF, ANY OF, and N OF M.

[0038] Preferably, the Authentication Source Adapters are pluggable modules configured to return standardized result objects comprising a source identifier, a confidence score, a signed evidence payload, and a verdict of pass, fail, or inconclusive.

[0039] Preferably, the adapters allow the system to ingest trust signals from multiple competing standards, including C2PA manifest chains, PRNU sensor fingerprint analysis, codec forensics, AI generative classifiers, and studio attestations.

[0040] Preferably, the cryptographic hashes are computed from the raw binary data of individual codec-encoded units sitting in the MDAT section of the media container

[0041] Preferably, the subset of binary samples for hashing is chosen according to a selection algorithm that may be deterministic, pseudorandom, adaptive, or biased toward specific frame types such as I-frames.

[0042] Preferably, the sample identity manifest is a compound record comprising the file identifier, selected sample positions, binary hashes, and a signed authentication provenance bundle recording the full chain of verifiers and policy conditions.

[0043] Preferably, assembly-time verification occurs in real time as the virtual video engine requests individual binary samples from remote media storage.

[0044] Preferably, the system computes the hash of each received sample and compares it against the registered manifest, enabling the system to block playback and signal unverified samples via a user interface.

[0045] Preferably, the authentication method is designed to survive re-encoding and other standard media transformations because the sample identity manifest and its provenance bundle exist outside the media file in an immutable registry.

[0046] Preferably, content authenticity verification is enforced as a mandatory condition for the validation of cryptographic video tokens that govern the conditions under which virtual video assembly is permitted.

[0047] Preferably, the system generates a traceable provenance chain report that surfaces the identities of contributing authentication sources.

[0048] The foregoing general description of the illustrative embodiments and the following detailed description thereof are merely exemplary aspects of the teachings of this disclosure and are not restrictive.

BRIEF DESCRIPTION OF THE FIGURES

[0049] Embodiments of the present disclosure will be discussed with reference to the accompanying figures wherein:

Figure 1 illustrates the hierarchical relationship between virtual video containers, distributed ledger governance, right to resolve references, and right to identify references;

Figure 2 depicts the operational workflow of the invention, tracing a media asset from initial human content creation through the authentication gate to verified, authenticated playback;

Figure 3 illustrates the universal adapter layer architecture, depicting how disparate authentication sources plug into a central policy engine;

Figure 4 depicts a sequence diagram for real-time assembly-time hash verification;

Figure 5 illustrates the structure of a Registered SIM Entry, comprising a sample identity manifest and a signed Authentication Provenance Bundle anchored to an immutable registry;

Figure 6 illustrates the process of generating a Sample Identity Manifest (SIM); and

Figure 7 depicts the stages of ingestion, adapter fan-out, result aggregation, and Boolean policy evaluation of the Authentication Policy Engine.

DESCRIPTION OF THE INVENTION

[0050] The reference to any prior art in this specification is not, and should not be taken as, an acknowledgement or any form of suggestion that such prior art forms part of the common general knowledge.

[0051] It will be understood that the terms “comprise” and “include” and any of their derivatives (e.g. comprises, comprising, includes, including) as used in this specification, and the claims that follow, is to be taken to be inclusive of features to which the term refers, and is not meant to exclude the presence of any additional features unless otherwise stated or implied.

[0052] In some cases, a single embodiment may, for succinctness and/or to assist in understanding the scope of the disclosure, combine multiple features. It is to be understood that in such a case, these

multiple features may be provided separately (in separate embodiments), or in any other suitable combination. Alternatively, where separate features are described in separate embodiments, these separate features may be combined into a single embodiment unless otherwise stated or implied. This also applies to the claims which can be recombined in any combination. That is a claim may be amended to include a feature defined in any other claim. Further a phrase referring to “at least one of” a list of items refers to any combination of those items, including single members. As an example, “at least one of: a, b, or c” is intended to cover: a, b, c, a-b, a-c, b-c, and a-b-c.

[0053] The term Right to Identify References refers to the cryptographic principle that digital video content must be authenticated as human-created at the binary sample level prior to being registered in a trust system and subsequently verified during real-time assembly.

[0054] The term Sample Identity Manifest (SIM) refers to a compound cryptographic record encapsulating both Sample Identity Data (binary hashes) and an Authentication Provenance Bundle (the audit record).

[0055] The term authentication policy engine means a configurable rules layer that acts as a gatekeeper for SIM registration.

[0056] The term Authentication Source Adapter refers to a discrete, pluggable software module designed to evaluate media against a specific forensic or cryptographic standard.

[0057] The term Virtualization Engine means the component responsible for interrogating media structure via a Sample Indexing Module to enable sample-level addressing and separation of container metadata from binary media data.

[0058] The term Video Token refers to a cryptographically validated control layer that governs assembly conditions.

[0059] The term Immutable Registry means a timestamped data store implemented via a distributed ledger or blockchain.

[0060] The term structural atoms means is a container unit (or box) in a MP4 file used to organize and store multimedia data and metadata inside the file. In the current terminology, an atom defines the structure of the file by separating information into logical sections such as file type information movie metadata audio/video track details.

[0061] The present invention represents a foundational infrastructure layer for digital media authentication and forensic security.

[0062] The invention addresses the global crisis of media and deepfakes by establishing a cryptographic record that proves whether digital content is authentic, human-created media at the individual binary sample level.

[0063] Unlike existing detection or provenance technologies that treat video as a sealed, monolithic file, the present invention operates as an independent reference layer beneath existing tools and delivery systems. It provides a universal adapter architecture that unifies disparate authentication methods such as C2PA, PRNU sensor fingerprinting, and AI detector ensembles into a single, policy-gated immutable record.

[0064] The system consists of several modules. The Authentication Policy Engine serves as the primary gatekeeper that loads signed authentication requirements to evaluate inbound media before it is permitted to enter the trust registry.

[0065] The system utilizes plurality of Authentication Source Adapters, which are pluggable modules that evaluate media against specific forensic or cryptographic standards. These adapters return standardized results including a source identifier, a confidence score, a signed evidence payload, and a verdict of pass, fail, or inconclusive.

[0066] A Sample Indexing Module works in conjunction with a virtualization engine to deconstruct media containers. This module interrogates structural atoms (such as the MOOV atom in MP4 files) to extract sample tables, timing metadata, and chunk offsets, enabling the system to address individual binary samples.

[0067] The Sample Identity Manifest (SIM) Registry stores the compound identity records that bind cryptographic sample hashes to an Authentication Provenance Bundle. This registry is preferably implemented as a distributed ledger to ensure the records are immutable and timestamped.

[0068] For end-user delivery, the system includes a Token Validation Module and an Assembly-Time Verification Module. The token module governs the conditions of assembly, while the verification module performs real-time hash comparisons of received binary samples against the registered SIM to confirm authenticity during playback.

[0069] The method initiates at an Ingestion Stage, where footage arrives with a reference to a signed Authentication Policy. This policy is a signed data object, rather than hardcoded logic, authored by rights holders or regulators to define the specific trust thresholds required for registration.

[0070] The Authentication Policy Engine fans out to required adapters and the results undergo a Boolean Policy Evaluation using composable operators such as ALL OF, ANY OF, and N OF M0. This allows for nested, graduated trust requirements, such as requiring both C2PA and a device registry check, or at least three forensic checks from a pool of four.

[0071] If an authentication source is unavailable or returns an inconclusive result, the engine may automatically fail the ingestion, skip the check based on policy redundancy, or route the asset to a human review interface.

[0072] Once the policy requirements are satisfied, the method proceeds to Sample Level Indexing step.

[0073] The virtualization engine extracts complete structural metadata (e.g., STBL, STTS, STCO) to map every discrete codec-encoded unit of media by its byte offset, length, and temporal position

[0074] Following indexing, the method initiates Cryptographic Identity Record Generation. A selection algorithm which may be deterministic, pseudorandom, or biased toward I-frames.

[0075] The method further computes cryptographic hashes of the raw binary data for each selected sample. These hashes are derived from the byte stream in the MDAT section of the file rather than the decoded pixel output, ensuring the verification is based on immutable binary identity rather than perceptual analysis.

[0076] These individual hashes are compiled into the Sample Identity Manifest (SIM), which is then registered alongside the Authentication Provenance Bundle. This compound record links the what (the binary hashes) with the how (the full record of contributing verifiers, policy versions, and confidence scores) .

[0077] The final stage of the method is Assembly-Time Verification that occurs when a user requests playback, retrieves the registered hashes from the SIM Registry and fetches binary samples from the remote source.

[0078] As the virtual video engine processes the stream, it computes the cryptographic hash of each received sample in real-time. These computed values are compared against the registered SIM hashes to confirm that the samples are identical to those verified at ingestion.

[0079] Furthermore, the system responds based on a pre-selected Verification Mode that includes a Strict Mode that blocks playback upon any mismatch; Advisory Mode that signals unverified samples to the user interface; and Audit Mode that logs all results for background compliance.

[0080] Every verification event is inherently traceable through the Provenance Chain Reporting stage The system can report exactly which samples were verified, under what token conditions the assembly occurred, and which authentication sources provided the original trust signals.

[0081] A primary advantage of the invention is that it is not an arms race. While generative AI models can produce visually convincing pixels, they cannot replicate the exact byte-level binary encoding of a real, camera-captured sample without access to that specific binary data.

[0082] The system is characterized by high computational efficiency and the heavy cost of forensic authentication is paid only once at the point of registration, while every subsequent playback inherits that trust through computationally trivial hash comparisons.

[0083] Furthermore, the invention survives re-encoding. The Universal Adapter Strategy ensures the architecture is future-proof. As the industry develops new forensic techniques or provenance standards, they plug into the ION system as new adapters without requiring changes to the core media pipeline.

[0084] Figure 1 illustrates the hierarchical relationship 100 between virtual video containers, distributed ledger governance, right to resolve references, and right to identify references.

[0085] The stack begins at the base with Virtual Video Containers 101, which establishes the Structure layer. This layer is responsible for separating the structural container of a video file from the media binary sample data, enabling sample-level addressing and efficient reference-based playback.

[0086] Above the structure layer sits the Distributed Ledger Governance 102, which serves as the Record layer. This layer utilizes distributed ledger technology to maintain an immutable, consensus-validated record of every content access transaction, ensuring transparency and security for all stakeholders.

[0087] The Right to Resolve References 103 functions as the Govern layer. It introduces a cryptographic video token that controls who can assemble content and under what specific conditions, such as licensing terms and execution constraints, at the moment of playback.

[0088] Finally, the apex of the stack is the Right to Identify References 104, which provides the Authenticate layer. This layer authenticates content as human-created at the individual binary sample level, ensuring that any assembly governed by the Resolve layer is verified as authentic media.

[0089] Figure 2 provides a schematic view of the end-to-end journey of a media asset 200, tracing it from initial creation to final delivered playback.

[0090] This workflow integrates all four layers (virtual video containers, distributed ledger governance, right to resolve references, and right to identify references) into a single, cohesive operational sequence.

[0091] The process initiates with content creation 201 that represents human-created video captured on physical hardware. This ensures the starting point is a genuine media asset before any digital processing or synthetic manipulation can occur.

[0092] Following creation 201, the media is subjected to the authentication gate 202, a multi-source policy evaluation gate. Thereafter, the Authentication Policy Engine evaluates the footage against a plurality of adapters, such as C2PA, PRNU, and AI Detectors, to determine its authenticity.

[0093] If the content passes the gate, it moves to Virtualisation and SIM Registration stage 203. During this stage 203, sample hashes and an authentication provenance bundle are generated and anchored to an immutable registry, locking the established trust into the structural data of the media.

[0094] The resulting Virtual Video Container 204 is distributed to end-users.

[0095] This container 204 contains no actual media data but is governed by a Video Token from The Right to Resolve References, which dictates the rules for how the video can be resolved and assembled.

[0096] During playback, the system performs Assembly-Time Verification 205 . This involves a real-time cryptographic hash comparison for every binary sample as it is fetched from remote storage, confirming that no modifications have occurred since registration.

[0097] Successful verification concludes in Authenticated Playback 206. At this final stage 206, every sample is verified, and the full provenance chain is traceable, providing the user with complete confidence in the content's integrity.

[0098] Figure 3 details the Universal Adapter Layer 300, which serves as the central unification point for diverse authentication technologies.

[0099] This layer 300 allows system architecture to ingest and orchestrate signals from competing and emerging standards without core changes.

[00100] This layer accepts pluggable inputs from multiple sources, including C2PA/Content Credentials 301 for validating manifest chains and PRNU/Sensor Fingerprints 302 for verifying camera sensor authenticity. These provide historical and hardware-based trust signals, respectively.

[00101] Furthermore, the layer integrates codec forensics 303 for analysing compression artifacts and AI Detectors 304 for generative media classification. These forensic tools identify technical inconsistencies that might suggest synthetic generation or splicing.

[00102] Additionally, it consumes studio attestations 305 for production chain-of-custody verification and capture device registries 306 for cross-referencing hardware certificates. This multi-signal approach ensures a robust, composable trust record.

[00103] The final output of this unified evaluation is stored in the Immutable SIM Registry 307. This registry 307 holds both the individual sample hashes and the comprehensive authentication provenance bundle, providing a permanent anchor for the asset's verified status.

[00104] Figure 4 depicts the sequence for Real-time Assembly-Time Hash Verification 400 during a playback session. This process ensures the integrity of the media at the moment it is de-virtualized and presented to the user.

[00105] When a user requests video playback 401, the system accesses the Virtual Video Container 402. This container 402 contains no media data, acting instead as a set of references and control parameters.

[00106] The Virtual Video Engine 403 validates the Video Token and initiates assembly. It begins fetching binary samples from various source files 404 while simultaneously retrieving registered hashes from the SIM Registry 405.

[00107] A Hash Comparison 406 is performed in real time for every sample fetched. Samples where the computed hash matches the registered hash (e.g., "a7f3...") are termed as VERIFIED, while mismatches (e.g., "ff02...") are flagged as UNVERIFIED.

[00108] The Verification Mode 407 determines the final system response to the comparison results. The system offers Strict Mode 407 to block playback on failure, Advisory Mode 407 to signal unverified samples to the user, or Audit Mode 407 for background logging.

[00109] Figure 5 illustrates the detailed structure of a Registered SIM Entry 500, which functions as a compound cryptographic record bound to the Immutable Registry 503.

[00110] This structure 500 ensures that the authentication record is both comprehensive and tamper-proof.

[00111] The entry consists of a Sample Identity component 501, which includes an ordered list of individual cryptographic hashes and a composite Root Hash. It also records metadata about the Selection Algorithm (e.g., I-frame biased) and total sample counts.

[00112] This is inextricably bound to an Authentication Provenance Bundle 502, which details the specific Policy applied and the individual results from contributing adapters.

[00113] Results display PASS/FAIL or NEGATIVE status for sources like C2PA, PRNU, and AI Detectors.

[00114] The bundle 502 also captures the evaluation timestamp and signed verifier identities. This compound architecture ensures the trust record is timestamped, anchored, and permanent, proving not just that a sample is a sample, but how it was verified.

[00115] Figure 6 is an illustration of the deconstruction of a Rendered Video File (MP4) 600 to generate the identity manifest. This file 600 is critical for establishing the ground truth of the human-created content.

[00116] The system interrogates the file structure, focusing on the MOOV (Structure) atom 601. It extracts critical tables including STBL, STTS, STSC, STCO, and STSZ, which provide the necessary map to locate media chunks in the container.

[00117] Using these structural records, the system maps individual codec-encoded samples within the MDAT (Media Data) section 602.

[00118] For each selected sample, the system computes a cryptographic hash 603 of raw binary data (e.g., "a7f3...").

[00119] These individual hashes 603, along with the file identifier, selection algorithm parameters, and root hash are compiled and signed to create the Sample Identity Manifest (SIM) 604 .

[00120] Figure 7 illustrates the Authentication Policy Engine 700, showing the step-by-step procedural flow from media intake to final registration. This represents the core logic that gates content into the ION trust system.

[00121] At the ingestion step 701, the footage is submitted along with a policy reference.

[00122] The Authentication Policy Engine 702 loads the signed requirements and fans out the footage to required Adapters such as C2PA 703, PRNU 703, and Codec Forensics 703, in parallel.

[00123] In the Result Aggregation stage 704, each adapter returns its standardized result object, including source ID and confidence score. These results then undergo Boolean Policy Evaluation 705 using operators like ALL OF, ANY OF, or N OF M.

[00124] Depending on the evaluation outcome, the system makes a gate decision. This results either in a Reject response 706 with a structured failure explanation or a command to Proceed to SIM 707 for successful assets.

[00125] Successful candidates conclude at SIM Registration 708. The sample hashes and the provenance bundle are inextricably anchored to an immutable registry, completing the authentication process.

[00126] In an alternative embodiment, the system architecture may be integrated directly into the hardware of a digital capture device to enable authentication at the point of origin. By operating the Authentication Policy Engine and specific adapters locally, the system interrogates the internal media buffer and MOOV atom to compute cryptographic hashes before file finalization.

[00127] A signed Sample Identity Manifest and Provenance Bundle may be transmitted to the Immutable Registry, establishing the Right to Identify References at the moment of creation and eliminating any gap between capture and ingestion.

[00128] It will be appreciated by those skilled in the art that the disclosure is not restricted in its use to the particular application or applications described. Neither is the present disclosure restricted in its preferred embodiment with regard to the particular elements and/or features described or depicted herein. It will be appreciated that the disclosure is not limited to the embodiment or embodiments disclosed, but is capable of numerous rearrangements, modifications and substitutions without departing from the scope as set forth and defined by the following claims.

Please note that the following claims are provisional claims only, and are provided as examples of possible claims and are not intended to limit the scope of what may be claimed in any future patent applications based on the present application. Integers may be added to or omitted from the example claims at a later date so as to further define or re-define the scope of the invention.