AUDIT SENTINEL AI

Audit Methodology White Paper

A Compliance-Defensibility Framework for AI-Assisted E/M and ICD-10 Billing Audits

Three-Pass AI Audit Architecture | Google Cloud Vertex AI

2023 AMA Revised E/M Standards | Current through the 2026 CPT Code Set

Document Version 1.0

Published: April 2026

Audit Sentinel, Inc. | auditsentinel.ai

Document Purpose and Intended Audience

This white paper documents the technical architecture, clinical coding methodology, and grading logic of the Audit Sentinel AI audit engine. It is intended as a reference document for compliance officers, internal and external auditors, health information management (HIM) professionals, billing and coding managers, healthcare attorneys, and any qualified reviewer who has been asked to evaluate or defend an Audit Sentinel audit output.

The methodology described herein is the basis on which Audit Sentinel produces its numeric and letter grades, deduction line items, compliance flags, and narrative findings. Where a provider, payer, or reviewer challenges an audit output, this document should be read in conjunction with the specific audit report in question. The grading rules, deduction values, and coding guideline references in this document are definitive; any apparent conflict between this document and marketing collateral should be resolved in favor of this document.

Audit Sentinel is an AI-assisted audit tool. It is not a certified coder, not a licensed attorney, and not a payer. Its outputs are educational and advisory in nature and are designed to support, not replace, the judgment of qualified human professionals. See Section 7 for the full limitations and disclaimer.

1. Executive Summary

Audit Sentinel AI is a clinical documentation and billing audit engine that evaluates Evaluation and Management (E/M) coding and ICD-10 diagnosis coding against current Current Procedural Terminology (CPT) guidelines. It is designed to support pre-bill and retrospective audit workflows for physician practices, hospitals, health information management departments, independent auditors, and compliance firms.

The engine is built on a three-pass pipeline running on Google Cloud Vertex AI under a Business Associate Agreement (BAA). The three passes are architecturally and functionally separated to ensure that (a) protected health information (PHI) is removed from the clinical note before any substantive analysis occurs; (b) the coding analysis is performed against a de-identified note only; and (c) the grading and scoring of the provider's submitted codes is produced by a distinct comparator pass with explicit, deterministic deduction rules. The pipeline uses a fast, lightweight frontier model for PHI redaction (Pass 1) and a high-capability frontier reasoning model for clinical audit and grading (Passes 2 and 3). Specific model versions in production are updated as improved models are released by the platform provider; current model versions are documented in the Audit Sentinel security portal and are available to customers under NDA.

The coding guideline basis is the 2023 American Medical Association (AMA) revised E/M standards, which extended the 2021 office-visit Medical Decision Making (MDM) and time-based framework to all E/M settings, together with CPT updates for 2024, 2025, and 2026. This includes the "must be met or exceeded" rule for time-based E/M selection and the revised telehealth code set.

Key Design Principles

- Privacy-by-architecture. PHI is redacted by an isolated upstream pass using the HIPAA Safe Harbor method before any clinical reasoning occurs.
- Guideline fidelity. MDM is evaluated on the 2023 AMA framework (Problems, Data, Risk) and time is evaluated on total practitioner time on the date of the encounter under the "met or exceeded" rule.
- Deterministic grading. The grade is computed from a fixed deduction table and a fixed letter-grade scale. The AI does not assign the final grade by discretion; it identifies findings that map to pre-defined deductions.
- Explainability. Every grade deduction is accompanied by a structured reason code, a human-readable narrative, and a reference to the specific element of the note or code set that triggered it.
- Defensibility. The audit output is educational and advisory; it is not a payer determination, a legal opinion, or a substitute for a qualified human auditor.

2. Purpose and Scope

2.1 What Audit Sentinel Audits

Audit Sentinel accepts a clinical note and, optionally, the provider's submitted billing codes for the corresponding encounter. The engine reviews:

- E/M level selection across outpatient office visits (CPT 99202–99205 and 99212–99215), hospital inpatient and observation care, emergency department services, nursing facility services, home or residence services, and office/outpatient consultations, where supported by the governing guideline set.
- Time-based E/M selection where the provider has documented total practitioner time on the date of the encounter and where time is a valid basis for code selection under current CPT rules.
- ICD-10-CM diagnosis code selection for specificity, clinical support in the documentation, and linkage to the E/M service for medical necessity.
- Modifier application, including required modifiers, unsupported modifiers, and incorrectly applied modifiers for the code set under review.
- Add-on codes and correctly bundled versus unbundled procedure combinations under National Correct Coding Initiative (CCI) edits.
- Documentation adequacy to support the level of service billed, including history and examination as clinically indicated, MDM elements, and time statements where applicable.

2.2 What Audit Sentinel Does Not Do

The following are explicitly outside the scope of the audit engine:

- **Legal advice.** Audit Sentinel does not provide, and its outputs do not constitute, legal advice. No attorney-client relationship is created by use of the engine.
- **Payer determinations.** Audit Sentinel does not determine whether a claim will be paid, denied, or recouped by any payer, including Medicare, Medicaid, or commercial plans. Payer policy varies and may differ from the CPT guideline basis used by the engine.
- **Fraud adjudication.** Audit Sentinel identifies coding variances and risk indicators, including potential over-coding. It does not determine intent and does not adjudicate fraud, abuse, or willful misconduct. Those determinations are reserved to qualified investigators and authorized authorities.
- **Clinical diagnosis or treatment advice.** The engine does not diagnose patients, recommend treatment, or second-guess clinical judgment. It evaluates whether the documentation as written supports the codes as submitted.
- **Non-E/M professional services outside declared scope.** Surgical global-period analysis, anesthesia time units, radiology component analysis, and pathology coding are out of scope except where they appear as modifiers or add-ons attached to an in-scope E/M encounter.

2.3 Intended Use

Audit Sentinel is intended as a high-confidence first-pass review that accelerates human auditor workflow, supports internal compliance monitoring, and surfaces documentation gaps before claims are submitted. Final coding and billing decisions remain the responsibility of the provider, the provider's coding staff, and any certified professional coder, health information professional, or compliance officer assigned to the encounter.

3. PHI De-identification Architecture

3.1 Pass 1: PHI Scrubber Overview

Pass 1 of the Audit Sentinel pipeline is a dedicated PHI scrubbing stage that executes on Google Cloud Vertex AI using a fast, lightweight frontier language model selected for latency and recall on entity-recognition tasks. Pass 1 is the first and only pass that touches the raw clinical note as submitted. Its sole function is to identify and redact protected health information per the HIPAA Safe Harbor de-identification method at 45 CFR § 164.514(b)(2), and to emit a de-identified note that is used as the input to all subsequent passes. Current model versions and revision history are maintained in the Audit Sentinel security portal and are available to customers under NDA.

No substantive clinical reasoning, coding analysis, or grading is performed in Pass 1. No downstream pass has access to the raw note. This architectural separation is a core compliance property of the system.

3.2 Safe Harbor Identifiers Redacted

Pass 1 identifies and redacts all eighteen Safe Harbor identifier categories. Each identifier is replaced with a standardized placeholder token so that downstream passes can still recognize that an entity existed in the original note without being exposed to its value.

#	Identifier Category	Standardized Placeholder
1	Names (patient, relatives, employers, household members)	[REDACTED_NAME]
2	All geographic subdivisions smaller than a state (street, city, county, precinct)	[REDACTED_ADDRESS]
3	ZIP codes (all digits, per Safe Harbor population thresholds)	[REDACTED_ZIP]
4	All elements of dates (except year) directly related to an individual, and all ages over 89	[REDACTED_DATE]

#	Identifier Category	Standardized Placeholder
5	Telephone numbers	[REDACTED_PHONE]
6	Fax numbers	[REDACTED_FAX]
7	Email addresses	[REDACTED_EMAIL]
8	Social Security numbers	[REDACTED_SSN]
9	Medical record numbers	[REDACTED_MRN]
10	Health plan beneficiary numbers	[REDACTED_PLAN]
11	Account numbers	[REDACTED_ACCT]
12	Certificate or license numbers	[REDACTED_LICENSE]
13	Vehicle identifiers and serial numbers, including license plates	[REDACTED_VEHICLE]
14	Device identifiers and serial numbers	[REDACTED_DEVICE]
15	Web URLs	[REDACTED_URL]
16	Internet Protocol (IP) addresses	[REDACTED_IP]
17	Biometric identifiers (finger, voice prints)	[REDACTED_BIOMETRIC]
18	Full-face photographs and any comparable images; any other unique identifying number, characteristic, or code	[REDACTED_OTHER]

3.3 User Responsibility for Pre-Submission Hygiene

The PHI scrubber is designed to be conservative: when in doubt, redact. However, automated Safe Harbor redaction cannot perfectly distinguish every patient-specific token from routine clinical vocabulary in every context. For example, a rare surname may overlap with a medical eponym, or a free-text string may contain an identifier that a reasonable reader would recognize only by context.

Customers are therefore instructed not to paste actual patient names, real medical record numbers, full street addresses, or other direct identifiers into the submission field. The scrubber is a defense-in-depth layer; it is not a license to submit raw identifiers. This guidance is surfaced in the product UI and in customer documentation.

3.4 What Is and Is Not Persisted from Pass 1

After Pass 1 completes:

- **Persisted:** the de-identified note only. This de-identified artifact is the clinical input used by Pass 2 and is stored as part of the audit record.
- **Not persisted:** the raw submitted note. The raw note is held only in volatile memory for the duration of Pass 1 execution and is discarded once the de-identified note is emitted. Raw PHI is not written to any durable store, is not available to Passes 2 or 3, and is not available for subsequent retrieval by any user, administrator, or support engineer.
- **Not used for model training.** Customer submissions, including the de-identified note and the raw note during its transient presence in Pass 1, are not used to train, fine-tune, or otherwise update the underlying foundation models.

4. E/M Coding Methodology

4.1 Guideline Basis

Audit Sentinel evaluates E/M coding against the 2023 AMA revised E/M guidelines, which extended the MDM-or-time framework introduced for office visits in 2021 to substantially all other E/M settings, including hospital inpatient and observation, emergency department, nursing facility, home or residence services, and consultations. The engine incorporates CPT updates for 2024, 2025, and 2026, including:

- The “must be met or exceeded” rule for time-based code selection. A time-based E/M code requires that the documented total practitioner time on the date of the encounter equals or exceeds the lower time threshold of the selected code.
- Revised and expanded telehealth E/M codes in accordance with CPT updates effective in 2025 and carried forward through 2026, including audio-only and audio-video distinctions where specified.
- Continued use of the three MDM elements: (i) Number and Complexity of Problems Addressed, (ii) Amount and/or Complexity of Data to be Reviewed and Analyzed, and (iii) Risk of Complications and/or Morbidity or Mortality of Patient Management.
- MDM level is set by the two-out-of-three rule across the three MDM elements.

4.2 MDM Framework (Pass 2)

Pass 2 of the pipeline, which runs on Google Cloud Vertex AI using a high-capability frontier reasoning model, performs a complete MDM-based analysis of the de-identified note. For each encounter, the engine evaluates each of the three MDM elements against the AMA MDM grid and assigns an element-level classification of Straightforward, Low, Moderate, or High. The overall MDM level is the level met or exceeded by at least two of the three elements.

MDM Level	Problems Addressed	Data Reviewed and Analyzed	Risk of Management
Straightforward	Minimal: 1 self-limited or minor problem	Minimal or none	Minimal risk of morbidity
Low	Low: 2+ self-limited; or 1 stable chronic; or 1 acute uncomplicated	Limited (one category met)	Low risk of morbidity
Moderate	Moderate: 1+ chronic with exacerbation; or 2+ stable chronic; or undiagnosed new problem with uncertain prognosis; or acute illness with systemic symptoms; or acute complicated injury	Moderate (one category met, with specified combinations of tests/notes/independent interpretation)	Moderate risk, including prescription drug management, decision regarding minor surgery with risk factors, social determinants significantly limiting diagnosis or treatment
High	High: 1+ chronic with severe exacerbation or threat to life/bodily function; or acute/chronic illness posing threat to life or bodily function	Extensive (one category met, including independent interpretation or discussion with external source)	High risk, including drug therapy requiring intensive monitoring for toxicity, decision regarding elective major surgery with identified risk factors, decision regarding emergency major surgery, decision regarding hospitalization, decision regarding DNR or to de-escalate care

The above is a structural summary of the AMA MDM grid and is not a substitute for the full AMA text. Pass 2 applies the complete grid, including category-specific combinations required to meet Data (Categories 1, 2, and 3).

4.3 Time-Based Code Selection

Where the encounter type permits selection by time and the note documents total practitioner time on the date of the encounter, Pass 2 also evaluates the time-based path and determines the highest defensible code between MDM and time. Under current CPT rules, the documented time must be met or exceeded for the lower bound of the selected code. If a time-based code is submitted but the note does not adequately document total time on the date of the encounter, a time-documentation deduction applies (see Section 6).

Pass 2 distinguishes total practitioner time on the date of the encounter from narrower historical constructs such as face-to-face time or counseling-dominated time. Only qualifying activities under current CPT definitions are considered, and only time personally spent by the reporting practitioner is counted.

4.4 Encounter Type Coverage

Pass 2 applies the appropriate code family and guideline subset for each of the following encounter types:

- **Office or other outpatient visits —** new patient (99202–99205) and established patient (99212–99215). MDM or total time on the date of the encounter.
- **Hospital inpatient and observation care —** initial, subsequent, and discharge services, evaluated under the 2023-extended MDM-or-time framework.
- **Emergency department services —** 99281–99285, evaluated by MDM. Time is not a basis for ED code selection.
- **Nursing facility services —** initial, subsequent, and discharge services, evaluated under the 2023-extended framework.
- **Home or residence services —** new and established patient codes, evaluated under the 2023-extended framework.
- **Consultations —** office/outpatient and inpatient consultations, where permitted and where the payer recognizes consultation codes.
- **Telehealth —** audio-video and audio-only codes per the CPT updates effective 2025 and carried into 2026, with appropriate modifier and place-of-service analysis.

5. ICD-10 and Medical Necessity Review

5.1 ICD-10-CM Validation

Pass 2 evaluates every submitted ICD-10-CM diagnosis code on four dimensions:

- **Validity.** The code exists in the current ICD-10-CM code set and is reported at the required level of specificity (e.g., required laterality, encounter type, or episode of care).
- **Clinical support.** The documentation in the de-identified note contains clinical findings, history, or assessment language that reasonably supports the diagnosis. Diagnoses stated only in a historical list or problem list without being addressed during the encounter are flagged as unsupported for this date of service.
- **Specificity.** Where ICD-10-CM provides a more specific code and the documentation supports it, the engine flags the use of an unspecified code as a specificity gap (e.g., I10 Essential hypertension may be appropriate, but a documented hypertensive emergency should be coded to a more specific category).

- **Sequencing.** The principal diagnosis or first-listed diagnosis is evaluated for consistency with the reason for the encounter as described in the note.

5.2 Medical Necessity Linkage

Medical necessity, for purposes of this audit, means that the documented diagnoses clinically justify the level of E/M service billed. Pass 2 evaluates whether the problems addressed, the data reviewed, and the risk of management — as reflected in the diagnoses and the note — are consistent with the submitted E/M level. A claim may be properly coded as to CPT selection yet be at risk on medical-necessity grounds if the diagnoses do not support the intensity of the service.

Where Pass 2 identifies a medical-necessity concern, it is reported as a distinct finding in the Pass 3 audit report. It does not by itself force a code change; it documents risk.

5.3 CCI, Bundling, and Add-On Codes

Pass 2 performs a procedural-edit review against the National Correct Coding Initiative (CCI) principles, including column 1 / column 2 edits and mutually exclusive edits, to identify unbundling risk. Add-on codes are checked for appropriateness and for the presence of a valid primary code. Missing or unsupported add-on codes and unbundled procedure pairs are reported to Pass 3 as discrete findings for deduction.

6. Grading Methodology

6.1 Pass 3: Billing Comparator and Grader

Pass 3 of the pipeline runs on Google Cloud Vertex AI using the same class of high-capability frontier reasoning model employed in Pass 2. It does not re-perform the coding analysis. Its function is to compare the provider's submitted codes against the ideal coding analysis produced by Pass 2 and to apply the deduction table below. Pass 3 outputs a structured JSON audit report that includes: a numeric grade (0–100), a letter grade, an itemized list of deductions with reason codes, a compliance flag (true/false), a risk profile, and a narrative findings summary.

The grading function is deterministic with respect to its findings: given the same set of findings, the deduction total and letter grade are fully determined by the table and scale in this section. The model's role in Pass 3 is to map findings to reason codes and produce the narrative; it does not assign the grade by discretion.

6.2 Deduction Table

Finding	Deduction	Compliance Flag
E/M perfect match (submitted = ideal)	**0 pts**	—

Finding	Deduction	Compliance Flag
E/M under-coded by 1 level	**-8 pts**	—
E/M under-coded by 2 or more levels	**-18 pts**	—
E/M over-coded by 1 level	**-20 pts**	High Risk
E/M over-coded by 2 or more levels	**-35 pts**	Critical — flag = true
Wrong code category (e.g., consult billed where office visit is correct)	**-25 pts**	—
Required modifier missing	**-10 pts**	—
Modifier present but unsupported by documentation	**-15 pts**	—
Modifier applied incorrectly	**-8 pts**	—
ICD-10 diagnosis missing or incorrect	**-10 pts**	—
Unbundled procedure (CCI edit violation)	**-15 pts**	—
Time-based code billed without adequate time documentation	**-12 pts**	—
Missing or unsupported add-on code	**-8 pts**	—

Scoring bounds. The grade is computed as 100 minus the sum of all applicable deductions, floored at 0 and ceilinged at 100. Deductions are additive; a single encounter may incur multiple deductions across distinct findings.

6.3 Compliance Flag Logic

The compliance flag is a boolean indicator (compliance_flag = true) that is asserted when the audit output reflects a finding of heightened regulatory exposure warranting immediate review by a compliance professional. The flag is asserted if any of the following conditions are met:

- E/M over-coded by 2 or more levels relative to the ideal analysis (Critical).
- A pattern of over-coded by 1 level findings may be elevated to a flag by the Pass 3 comparator when combined with other risk indicators such as unsupported modifiers, unbundling, or time-based codes without time documentation. The single-level over-code finding, standing alone, is a High Risk indicator but does not by itself assert the flag.

The compliance flag is intentionally conservative. It is not a determination of fraud or abuse; it is a signal that a qualified human reviewer should re-examine the encounter.

6.4 Grade Scale

Letter Grade	Numeric Range	Label	Interpretation
A	95 – 100	**Optimized**	Documentation and coding are aligned. No material deductions.
A-	90 – 94	**Compliant**	Minor, non-material issues. Claim is defensible as submitted.
B	80 – 89	**Minor Gap**	One or more documentation or coding gaps. Correct before submission where possible.
C	70 – 79	**Needs Improvement**	Multiple findings or a single material finding. Human review recommended.
D	60 – 69	**High Risk**	Material exposure, often including single-level over-coding or unsupported modifiers. Human review required.
F	0 – 59	**Critical**	Significant exposure, typically including two-or-more-level over-coding or multiple material findings. Compliance flag likely asserted.

6.5 Rationale for the Scale

The scale is calibrated so that a defensible, well-documented encounter reliably grades in the A / A- band; a claim with minor documentation gaps that do not change the billed code grades in the B band; and a claim with material over-coding exposure drops sharply. The deduction of 20 points for a single-level over-code and 35 points for a two-or-more-level over-code is deliberately steep: under-coding harms the provider's own revenue, but over-coding creates payer and regulatory exposure, and the grading reflects that asymmetry.

7. Limitations and Disclaimer

7.1 AI-Generated Output

Audit Sentinel outputs are generated in part by large language models running on Google Cloud Vertex AI. Like all large language model systems, the engine can make mistakes, including misreading clinical context, overlooking a subtle documentation element, misclassifying a diagnosis, or producing a narrative that is well-formed but factually incorrect in a specific particular. The three-pass architecture,

deterministic deduction rules, and structured JSON output are designed to reduce these risks and to make them auditable, but they do not eliminate them.

7.2 Not a Substitute for a Qualified Human Auditor

Audit Sentinel does not replace a Certified Professional Coder (CPC), Certified Coding Specialist (CCS), Registered Health Information Administrator (RHIA), Registered Health Information Technician (RHIT), certified compliance professional, or any other qualified human auditor. The engine is intended as a decision-support and first-pass review tool. Final coding decisions, billing decisions, compliance determinations, and any disclosure or self-reporting decisions must be made by qualified humans with full context, access to the raw record, and applicable policies.

7.3 Educational and Informational Use

Audit outputs, including numeric grades, letter grades, deductions, compliance flags, and narrative findings, are provided for educational and informational purposes. They are not:

- Legal advice. No attorney-client relationship is created by use of the engine.
- Payer determinations. Payer policy and local coverage determinations may differ from the CPT guideline basis used by the engine.
- Medical advice or a diagnosis of any patient.
- A determination of fraud, abuse, or intent.
- A certification, accreditation, or endorsement by the American Medical Association, the Centers for Medicare & Medicaid Services, or any other regulatory or standards body.

7.4 Guideline Updates

CPT and ICD-10-CM guidelines are updated at least annually. Audit Sentinel tracks published AMA and CMS updates and updates its guideline basis accordingly. At the time of publication of this white paper, the engine reflects the 2023 AMA revised E/M standards and CPT updates current through the 2026 code set. Customers reviewing historical audits should note the guideline version in effect at the time the audit was run; an encounter that was correctly coded under one year's guidelines may score differently under a later year's guidelines.

7.5 Use of Outputs in Proceedings

Audit Sentinel outputs may be shared with internal compliance personnel, external auditors, legal counsel, and, at the customer's discretion, payers or investigators. Customers retain full ownership of their audit outputs. Audit Sentinel makes no warranty, express or implied, that any specific audit output will be accepted as authoritative in any administrative, civil, or criminal proceeding. The engine is a tool; its outputs are evidence of the documentation-and-coding analysis it performed, not evidence of any ultimate legal or regulatory conclusion.

8. Data Handling

8.1 What Is Stored

For each audit, the following artifacts are persisted in the customer's tenant:

- The de-identified note emitted by Pass 1, with all eighteen Safe Harbor identifier categories replaced by standardized placeholders.
- The structured JSON audit report emitted by Pass 3, including the numeric grade, letter grade, itemized deductions with reason codes, compliance flag, risk profile, and narrative findings.
- The submitted billing codes (CPT/HCPCS, ICD-10-CM, modifiers, and add-ons) that were compared against the Pass 2 ideal analysis.
- A timestamp (date and time in UTC) for the audit, along with a tenant identifier and an audit identifier.
- Metadata describing the guideline version in effect at the time of the audit, so that historical audits remain interpretable after subsequent guideline updates.

8.2 What Is Not Stored

- The raw submitted note. Raw PHI is held only in volatile memory for the duration of Pass 1 execution and is discarded once the de-identified note is emitted.
- Any Safe Harbor identifier in its original form. Only standardized placeholders appear in persisted artifacts.
- Model weights, prompts, or any derivative of customer submissions used to train, fine-tune, or update the underlying foundation models. Customer submissions are not used for model training.
- External copies. Audit artifacts are stored in the customer's tenant within the Audit Sentinel platform; they are not forwarded to any third party other than Google Cloud as the sub-processor for Vertex AI model execution.

8.3 Retention Policy

Audit artifacts are retained for the period specified by the customer's subscription plan (by default, 30 to 90 days of in-product history, depending on plan, with longer retention available to higher-tier plans). Customers may export their audit artifacts as PDF, CSV, or encrypted archive at any time during the retention window. At the end of the retention period, or upon customer-initiated deletion, audit artifacts are removed from production storage on a rolling schedule and from backups in accordance with the Audit Sentinel backup lifecycle. Customers requiring extended retention for regulatory, litigation-hold, or self-audit purposes should export and preserve artifacts in their own record systems.

8.4 Security and Access Control

Data in transit is protected using TLS 1.2 or higher. Data at rest is encrypted using AES-256 or equivalent. Access to customer tenants is role-based; administrative access by Audit Sentinel personnel is logged and is limited to personnel with a documented business need. Audit Sentinel maintains a Business Associate Agreement (BAA) available to customers on eligible plans, and executes a corresponding BAA with Google Cloud for the Vertex AI sub-processing relationship.

8.5 Incident Response

In the event of a suspected security incident affecting customer data, Audit Sentinel will notify affected customers in accordance with the timelines and procedures set forth in the applicable BAA and in the Audit Sentinel Terms of Service. Customers are responsible for their own onward notification obligations, including those under HIPAA Breach Notification rules and applicable state law.

Appendix A. Structured Audit Report Output

Each audit produces a structured JSON payload whose top-level fields are summarized below. The full field set is documented in the Audit Sentinel API reference; this table describes the compliance-relevant core.

Field	Type	Description
audit_id	string	Unique identifier for the audit.
timestamp_utc	string (ISO 8601)	UTC timestamp at which the audit was performed.
guideline_version	string	CPT/ICD-10 guideline version in effect at the time of the audit (e.g., "CPT 2026").
submitted_codes	object	Provider-submitted CPT/HCPCS code, ICD-10-CM codes, modifiers, and add-on codes.
ideal_analysis	object	Pass 2 output: ideal E/M level, MDM breakdown, time analysis, ICD-10 review, medical-necessity findings.
deductions	array	Ordered list of deductions, each with reason_code, description, points, and supporting narrative.
numeric_grade	integer (0–100)	Final score after deductions, floored at 0 and ceilinged at 100.
letter_grade	string	One of: A, A-, B, C, D, F.
compliance_flag	boolean	True when the finding set asserts the compliance flag per Section 6.3.

Field	Type	Description
risk_profile	string	One of: Optimized, Compliant, Minor Gap, Needs Improvement, High Risk, Critical.
narrative_findings	string	Human-readable summary of the audit, produced by Pass 3.

Appendix B. Glossary of Key Terms

- **CPT.** Current Procedural Terminology. A code set maintained by the American Medical Association describing medical procedures and services.
- **E/M.** Evaluation and Management. The CPT code family used to report visits and other non-procedural professional services.
- **MDM.** Medical Decision Making. The AMA framework for selecting an E/M level based on Problems, Data, and Risk.
- **ICD-10-CM.** International Classification of Diseases, Tenth Revision, Clinical Modification. The U.S. diagnosis code set.
- **CCI.** National Correct Coding Initiative. CMS-published edits that identify improperly unbundled or mutually exclusive procedure pairs.
- **Safe Harbor.** The HIPAA de-identification method set forth at 45 CFR § 164.514(b)(2) that, when applied, renders information not individually identifiable for HIPAA purposes.
- **PHI.** Protected Health Information, as defined by HIPAA.
- **BAA.** Business Associate Agreement. The HIPAA-required contract between a covered entity (or another business associate) and a business associate.

Document Control

Field	Value
Document title	Audit Sentinel AI — Audit Methodology White Paper
Version	1.0
Effective date	April 2026
Guideline basis	2023 AMA revised E/M standards; CPT updates current through the 2026 code set
Owner	Audit Sentinel, Inc. — Compliance

Field	Value
Intended audience	Compliance officers, auditors, HIM professionals, healthcare attorneys, payers
Supersedes	N/A (initial version)