This model Standard Operating Procedure (SOP) establishes a legally defensible, enterprise-grade framework for integrating Artificial Intelligence (AI) into daily law enforcement operations.

It is structured specifically to align with federal Criminal Justice Information Services (CJIS) data security mandates, insulate your agency against liability, protect Personally Identifiable Information (PII), and withstand scrutiny from defense attorneys and internal auditors.

STANDARD OPERATING PROCEDURE: USE OF ARTIFICIAL INTELLIGENCE SYSTEMS

Subject: Use of Artificial Intelligence Systems	**Policy Number: 2026-ADM-042**
Effective Date: May 26, 2026	**Review Date:** Annually
Approved By: Office of the Chief of Police	**Classification:** Unclassified / Internal Policy

I. PURPOSE & SCOPE

The purpose of this policy is to establish clear operational guidelines and mandatory data security guardrails for the deployment and use of Artificial Intelligence (AI) technologies, specifically Large Language Models (LLMs), text generators, and automated narrative tools, within this Department.

This policy applies to all sworn officers, civilian personnel, contractors, and specialized units utilizing any department-authorized AI platform. It governs the generation of incident reports, investigative supplemental narratives, and warrant scaffolding.

II. DEFINITIONS

- **Criminal Justice Information (CJI):** All CJIS-regulated data necessary for law enforcement performance, including biometric data, identity history, biographic data, property data, and case/investigative narratives.

- **Personally Identifiable Information (PII):** Information that can be used to distinguish or trace an individual's identity, either alone or when combined with other personal or identifying information (e.g., SSN, date of birth, driver's license number, medical history).
- **Isolated Enterprise Environment:** A secure software architecture utilizing private cloud infrastructure (such as AWS Bedrock or an isolated VPC) where data inputs are contractually isolated, encrypted, and legally barred from being retained or used for third-party model training.
- **Public AI System:** Any commercial, consumer-facing AI utility (e.g., standard ChatGPT, Gemini, Claude) that retains user prompts to optimize or train public foundational models.

III. MANDATORY INFRASTRUCTURE & DATA PRIVACY GUARDRAILS

To prevent catastrophic privacy breaches, data leakage, and statutory violations, the department shall enforce absolute structural isolation of all data inputs.

1. Absolute Prohibition of Public AI Utilities

- **Directive:** Personnel are strictly prohibited from entering any operational data, CJI, PII, witness names, case numbers, or narrative bullet points into any public AI system or browser extension.
- **Reasoning:** Entering data into public models constitutes an unencrypted federal data leak and invalidates the secure chain of custody required for criminal prosecution.

2. Authorized Enterprise Architecture Only

All authorized AI utilities deployed within this Department must operate on a closed backend featuring:

- **Zero Training Retention:** A binding service-level agreement (SLA) contractually guarantees that no user inputs or generated outputs are stored, logged, or utilized by the vendor or cloud host for model optimization.
- **Data Segregation:** Database structures utilizing schema-level multi-tenancy to ensure this agency's operational data is physically or logically firewalled from all external entities.
- **Data Encryption:** Mandatory AES-256 encryption at rest and TLS 1.3 encryption in transit for all data flowing between the department interface and the AI API endpoints.

IV. DATA RESTRICTIONS & INPUT INGESTION STANDARDS

To minimize litigation risks and protect civilian privacy, personnel must strictly filter information before passing text to the AI engine.

1. Ingestion Restrictions for PII

When utilizing AI engines to draft narratives from field notes, personnel shall replace explicit core identifiers with generic placeholders during the generation phase wherever practical, unless operating within a fully audited, end-to-end encrypted local application.

- **Prohibited Ingestions:** Social Security Numbers (SSNs), juvenile identities, financial account numbers, unverified medical history records, and explicit victim contact data shall not be routed through external prompt pipelines.

2. Mandatory Pre-Classification Review

Before any investigative file, surveillance log, or witness statement is processed by an AI agent, the user must verify that the material does not contain classified federal intelligence or multi-jurisdictional non-disclosure data.

V. ACCOUNTABILITY & OPERATIONAL USE OF OUTPUTS

AI is an administrative tool; it lacks legal standing, police powers, and statutory authority. Ultimate accountability rests solely on the human operator.

1. The Human-in-the-Loop Imperative

- **Directive:** No AI-generated narrative, timeline, or warrant scaffolding may be exported, attached to a case file, or submitted to a supervisor or magistrate without thorough manual verification, editing, and explicit sign-off by the handling officer.
- **The Review Standard:** Personnel must line-read every sentence generated by the AI to verify absolute factual alignment with original field notes, physical evidence, and raw observation.

2. Elimination of Hallucinations and Subjective Language

Personnel shall ensure that the AI-generated text is free from technical errors or "hallucinations" (fabricated data). Narratives must strictly utilize objective, third-person, past-tense phrasing. Subjective speculation regarding a suspect’s intent, state of mind, or unverified emotional state must be manually expunged before submission.

VI. LITIGATION AND COURTROOM DEFENSE PREPARATION

To insulate the agency from discovery challenges, Brady/Giglio motions, and defense attacks on evidence integrity, the platform must maintain strict audit trails.

1. Immutable Activity Auditing

The department’s AI platform must automatically log and maintain an uneditable audit trail of all transactions. For every user session, the system shall record:

1. The unique user ID, date, and timestamp of the session.
2. The exact input data/notes provided by the officer.
3. The specific system prompt and foundational model used.
4. The exact unedited output generated by the AI.

Retention Policy: These immutable access logs shall be securely retained in an append-only database table for a period aligning with the department's statutory evidence retention mandates for the associated crime classification.

2. Subpoena Readiness & Transparency

- **Policy Position:** The utilization of an AI assistant to organize notes, fix grammar, or structure a standard narrative layout is an administrative drafting aid, equivalent to using a word processor or dictation software.
- **Discovery Strategy:** Because the platform records an immutable log of the raw input vs. the generated output, the agency can confidently defeat defense claims of "fabricated evidence" or "altered narratives" by providing the clean, verified data pipeline upon an official discovery request.

VII. DISCIPLINARY SANCTIONS

Violations of this policy, specifically the unauthorized entry of sensitive department investigations or PII into public AI tools, or the failure to review an AI-generated warrant that leads to a defective judicial submission, constitute severe administrative misconduct.

Such infractions will be subject to immediate internal investigation and disciplinary action, up to and including suspension and termination of employment, as well as potential civil or criminal exposure under federal privacy and data security statutes.