PUBLIC POLICY • CRIMINAL JUSTICE • ARTIFICIAL INTELLIGENCE

The Record Outlives the Decision

AI in Corrections and Sentencing: Nevada Safeguards

A Member's Policy Brief for the Nevada Sentencing Commission

Jarret A. Orcutt | Member, Nevada Sentencing Commission

Prepared October 29, 2026 for the Commission's November 20, 2026 meeting

Reflects public information as of October 28, 2026.

This brief presents one member's views and recommendations. It does not state an adopted position of the Nevada Sentencing Commission.

Summary

For AI in corrections and sentencing, a less visible risk lies in the record. An AI-generated narrative, summary, or assessment may become part of an official record that affects a person's liberty long after the document was created.

Artificial intelligence is reaching criminal justice agencies through procurement and new capabilities added to existing software. It can support transcription, translation, administrative work, and education. Police reports, presentence investigations, disciplinary records, classification assessments, supervision notes, and parole materials remain consequential when labeled administrative. An error introduced during drafting may be repeated, relied on, and difficult to reconstruct. The immediate policy question is whether agencies and affected people can identify material AI assistance, verify underlying facts, preserve relevant sources, and correct errors before those errors influence a decision.

Nevada has existing protections, including an executive-branch AI policy, judicial guidance, narrow statutory human-decision requirements, and procedures for challenging certain inaccurate records. None of them is a uniform AI-specific framework for the life of a criminal justice record. [1, 2, 3, 4, 5, 6]

This brief proposes **four record-integrity safeguards** for evaluation: meaningful human oversight; disclosure and preservation of AI-assisted records; accessible correction and challenge procedures; and auditability with confidentiality protections. It also proposes **a separate educational-data protection** to prevent learning activity in custody from being repurposed for disciplinary or parole risk profiling.

> **I recommend that the 2027 Legislature consider authorizing the Nevada Sentencing Commission to survey AI use in sentencing, corrections, supervision, and parole, including law enforcement reports that enter those records, and to publish best practices informed by its findings.** This brief is offered for the Commission's consideration without requesting an immediate vote, directive, or agency assignment.

1 The record outlives the decision

A disciplinary report may later inform classification or parole consideration. A probation note may influence supervision decisions. A presentence investigation may shape a judge's understanding of a defendant's history. When AI drafts, condenses, or reformulates information, a plausible sentence can

conceal an unsupported inference, omitted qualification, mistranscribed statement, or incorrect source attribution.

The Council on Criminal Justice's September 2026 police report-writing case study uses a hypothetical scenario to examine consequences for charging, negotiation, sentencing, accountability, and evidentiary challenges. [7]

The same reasoning extends to corrections. Consider a hypothetical example. A camera system flags two people exchanging a commissary item. An AI tool helps turn the flag into a disciplinary report. A sanction follows, and the report becomes part of the person's file. Years later, a parole panel reviews that file. This is an illustration, not a documented Nevada incident. At each stage, the question is the same: who can check or correct the record?

The operational danger is cumulative reliance: later readers may treat earlier narrative text as an independently verified fact. The origin and revision history of consequential records should remain inspectable as information moves between agencies.

Documented examples

Synthetic victim portrayal at sentencing. In *State v. Horcasitas*, the Arizona Court of Appeals affirmed a manslaughter conviction but vacated the sentence on September 30, 2026, after the sentencing judge relied on an unreliable AI-generated depiction of the deceased victim speaking. The court addressed the particular depiction's presentation of imagined thoughts as the victim's own. [8]

Fabricated authorities in parole proceedings. ABC reported in September 2026 that proceedings concerning Susan Neill-Fraser's parole conditions exposed nonexistent legal authorities in a Tasmanian Parole Board document associated with AI use. [9]

Inaccurate criminal filings in Nevada County, California. The *Sacramento Bee* reported in April 2026 that Nevada County prosecutors acknowledged AI-associated errors in four criminal cases. These cases illustrate the risk of generated legal content being used without adequate source checking. [10]

None of these happened in Nevada. They are reasons to look at record integrity here, not a measure of how often errors occur.

2 Identifiable human responsibility

An officer may spend less time preparing a first draft, a probation employee may find case information more quickly, and a student in custody may receive individualized practice and feedback. Those potential benefits depend on using tools appropriate to the task and verifying consequential output.

The proposed standard is human in the lead, not merely human in the loop: an accountable official must have the authority and responsibility to verify, correct, or reject consequential AI output. [11] An identifiable employee should understand the tool's limits, verify material facts against underlying sources, have authority to reject or correct output, and remain accountable for the official action. A signature without meaningful review is insufficient. Verification and review should increase with the potential effect on liberty, safety, housing, discipline, or release.

The Council on Criminal Justice's principles and decision framework, the Urban Institute's corrections analysis, and the Vera Institute's accountability principles offer relevant approaches to evaluation before adoption and oversight during use. They inform this proposal. [7, 12, 13, 14]

3 Four record-integrity safeguards

Safeguard 1 Human oversight with authority to reject

For a consequential document or recommendation, the responsible employee should independently review material factual assertions against available underlying sources. AI output should not supply the sole factual basis for a custodial disciplinary sanction, liberty-affecting classification decision, parole recommendation, or comparable adverse action. Review should be documented, and employees should have clear authority to reject or correct AI output they cannot independently verify.

Requirements should match the function. Transcription assistance, statistical risk instruments, and generative decision-support systems carry different risks. Review of a factual narrative and validation of a statistical instrument require different methods.

Safeguard 2 Disclosure and preservation of sources and drafts

When AI materially generates or revises an official narrative, the final document should identify the tool and describe its role. The responsible employee should verify factual content, and the initial machine-produced draft should be retained for the applicable retention period alongside the final report. Relevant source audio, video, records, or structured inputs needed to reconstruct material assertions should remain available under lawful evidence, privacy, and retention rules.

A definition of material assistance should cover substantive allegations, factual summaries, findings, and recommendations while distinguishing routine spelling or formatting corrections. Coverage should turn on function and consequence, so a product labeled an editor, transcription tool, or summarizer does not escape review when it materially shapes the record.

Safeguard 3 Meaningful correction and challenge procedures

A person whose sentencing, classification, discipline, or parole outcome may be materially affected should receive notice of AI assistance at an appropriate procedural stage and a practicable way to identify factual inaccuracies or material omissions. The process should provide a human reviewer independent of the disputed document's preparation, a written disposition, preservation of the dispute, and transmission of corrections to subsequent users.

Where an adverse decision is imminent, the process should provide expedited review and require the decision-maker to document how the dispute was handled, including whether reliance on disputed material should be deferred as permitted by law. Necessary language access, disability access, and access through counsel should be considered. The process should complement NRS 176.156, NAC 213.524, existing grievances, discovery rules, and constitutional remedies. Confidentiality and institutional safety protections should preserve a meaningful opportunity to contest the central facts. [4, 5]

Safeguard 4 Auditable systems and accountable reporting

Agencies should retain logs proportionate to the tool and consequence: user identity and role, date and purpose, relevant product or model version, material source references, consequential revisions or overrides, and access or disclosure events. Logs should make unauthorized alteration detectable. Aggregate annual reporting should identify systems, purposes, procurement and deployment status, significant incidents, audit findings, and corrective actions.

Public reporting should protect information about victims, incarcerated people, students, and security architecture. NRS 176.01385 protects specified information collected or stored by the Department of Sentencing Policy. [15]

California as a starting model

California's SB 524 added Penal Code §13663, effective January 1, 2026. For covered systems that automatically draft police-report narratives from specified audiovisual sources or an officer's dictated report, it requires identification and disclosure of AI use, a verifying signature, retention of the first AI-generated draft for the official report's retention period, and an audit trail identifying the user and relevant footage. It also distinguishes an AI draft from an officer's statement and restricts vendor data use, while permitting specified access for troubleshooting, bias mitigation, accuracy improvement, and system refinement. [16]

Nevada should evaluate extending calibrated protections to material AI-assisted correctional incident reports, classification summaries, supervision records, and parole recommendations. The review should address transcripts, translated content, and other source types.

Four questions about consequential records

The safeguards reduce to four questions every agency using AI on a consequential record should be able to answer:

1. Who verifies the record?
2. How are AI use and source information disclosed and retained?
3. How can material inaccuracies be challenged?
4. How are the systems audited?

4 A separate protection for educational AI data in custody

Bounded educational tools can support tutoring, literacy, vocational preparation, career exploration, and reentry planning. Educational engagement should not become an additional source of opaque surveillance merely because a digital tool was used.

Alongside the four record-integrity safeguards, this brief proposes a statutory or contractual restriction on repurposing educational AI prompts, tutoring exchanges, drafts, and academic learning histories for disciplinary classification, behavioral profiling, or parole risk scoring. The restriction should define covered systems and records, permit only narrowly defined exceptions required by applicable law or immediate safety needs, and require access logs and review of exceptional disclosures.

FERPA applies to eligible Prison Education Programs. Federal guidance explains that incarcerated students retain rights to inspect and seek amendment of education records, and that disclosure to correctional authorities generally requires consent unless an exception applies. Not every AI interaction is an education record. The proposed protection would address educational-use data beyond FERPA's coverage, with precise definitions and lawful exceptions. [17]

A fifth question follows: is educational data from custody being repurposed for discipline, classification, or parole?

5 Evaluation before adoption and procurement

Before purchasing or activating a high-impact AI function, an agency should assess its intended task, affected populations, alternatives, accuracy, failure modes, data retention, vendor access, cybersecurity, accessibility, and consequences of error. The assessment should identify the official accountable for approval and continued use.

Risk-scoring and classification systems should be validated on the populations they serve and assessed for material disparities using methods appropriate to the instrument.

Procurement terms should address independent audit and testing rights, data-use restrictions, breach notification, source preservation, access to logs, corrective action, and termination for material noncompliance. Evaluation should also identify staffing, storage, training, and vendor costs, including burdens on smaller agencies.

6 Nevada's starting point

Executive-branch policy. Nevada's published, signed *Policy on the Responsible and Ethical Use of Artificial Intelligence in Nevada State Government Executive Branch* establishes baseline expectations for transparency, security, accountability, and human oversight. It identifies official report generation and AI-generated insights for policy decisions as approval-required uses. Its executive-branch scope does not establish equivalent requirements for local governments or courts, or a statutory right for affected people to obtain a first draft and challenge its contents. [1]

Judicial guidance. The Administrative Office of the Courts' *AI Guide for Judicial Officers* addresses hallucinated citations, synthetic evidence, accuracy, privacy, and human judgment. It is guidance, not an enforceable statewide disclosure or correction rule. [6]

Human-decision precedents. NRS 414.099 requires specified emergency plans to ensure that final emergency-response planning and resource-allocation decisions are made by a natural person. NRS 704.1833 bars reliance solely on AI for specified emergency utility service reductions or shutoffs. These are narrow provisions, but they offer Nevada examples of expressly assigning responsibility to people. [2, 3]

Existing challenge procedures. NRS 176.156 affords parties an opportunity to object to factual errors in presentence or general investigation reports. Its specified post-sentencing correction mechanism permits a court-ordered correction when the prosecutor and defendant stipulate within 180 days after entry of the judgment of conviction. NAC 213.524 provides a conditional route to request reassessment of a parole risk level within 45 days of the relevant meeting. Neither provision was designed specifically for AI-assisted records or expressly requires disclosure of AI use. [4, 5]

Legislative history. Senator Dina Neal's SB 199 initially proposed prohibiting peace officers from using AI to generate investigative or police reports. Amendment No. 304, adopted April 21, 2025, replaced that prohibition with agency AI-use policy requirements, including reporting and facial-recognition uses. The bill was re-referred to Senate Finance and did not become law. [18, 19] For 2027, the Joint Interim Standing Committee on Government Affairs approved BDR –504 on August 27, 2026 to establish standards for state and local governmental AI use, including transparency, accountability, and protections against bias. It is an approved bill draft request, not an enacted law. [20, 21]

7 One recommendation for the 2027 session

Nevada is not starting from zero, but it still needs a focused picture of how AI tools are used and planned across sentencing, corrections, supervision, and parole, and of the records and decisions those tools affect. Existing law gives the Commission broad study and recommendation duties and requires the Department of Corrections and the Division of Parole and Probation to provide available statistical information or research and assist in compiling and developing information. [15, 22] Subpoenas follow separate procedures. [15] An express AI-specific legislative charge would clarify the survey's scope, participation by additional agencies, reporting expectations, and resource needs.

> **I recommend that the 2027 Legislature consider authorizing the Nevada Sentencing Commission to survey AI use in sentencing, corrections, supervision, and parole, including law enforcement reports that enter those records, and to publish best practices informed by its findings.**

We must understand these tools before we can lay the guardrails. A survey of this kind would begin with the Department of Corrections, the Division of Parole and Probation, the Parole Board, the courts, and relevant local agencies, and with the questions above. It should distinguish generative AI, machine learning, conventional automation, and tools whose functions remain unverified. It should hear from prosecutors, defense counsel, courts, correctional and supervision staff, parole authorities, incarcerated and formerly incarcerated people, educational institutions, privacy experts, and technology personnel.

The four record-integrity safeguards and the educational-data protection in this brief are offered as a starting point for that work, not as completed Commission findings.

Conclusion

The record outlives the decision. Nevada should determine what safeguards follow AI-assisted information through the criminal justice system and how people can identify and correct mistakes. That work starts with a clear charge to understand the tools: authority for the Commission to survey them and to publish best practices informed by its findings. Useful technology and reliable records both depend on identifiable human responsibility.

About the author

Jarret A. Orcutt serves on the Nevada Sentencing Commission, appointed by Governor Joe Lombardo in 2025. He is Coordinator of Educational Partnership Programs at Truckee Meadows Community College, which include the college's education programs for justice-involved students, and a Master of Public Administration candidate at the University of Nevada, Reno. He also teaches noncredit artificial intelligence courses for adults in the community through TMCC's EPIC program, including *Advanced AI: Workflows, GPTs, and Intelligent Agents*. He spent seventeen years in the custody of the Nevada Department of Corrections and returned to college at TMCC while still incarcerated. His March 2026 brief, *From Custody to Career: College as the Framework for Reentry*, argued that education is the framework for reentry, and his book *Redesigning Justice-Involved Education* is forthcoming from Taylor & Francis. The views in this brief are his own and do not represent Truckee Meadows Community College.

Contact: Jarret A. Orcutt, jorcutt@tmcc.edu

Suggested citation: Orcutt, J. A. (2026). *The Record Outlives the Decision. AI in Corrections and Sentencing: Nevada Safeguards*. Member's policy brief for the Nevada Sentencing Commission, prepared October 29, 2026.

References and source notes

1. Nevada Office of the Chief Information Officer. Policy on the Responsible and Ethical Use of Artificial Intelligence in Nevada State Government Executive Branch, signed policy linked from the state's governance page; approval-required uses at page 3 and human oversight at pages 3–4.
2. Nevada Legislature. NRS 414.099, human responsibility for specified emergency planning and resource-allocation decisions.
3. Nevada Legislature. NRS 704.1833, specified emergency utility service decisions and AI.
4. Nevada Legislature. NRS 176.156, disclosure, objections, and the specified post-sentencing correction procedure.
5. Nevada Legislature. NAC 213.524, reassessment of a parole risk level under specified conditions.
6. Nevada Administrative Office of the Courts. AI Guide for Judicial Officers, including accuracy, synthetic evidence, privacy, and human-judgment guidance.
7. Council on Criminal Justice. National Task Force Releases Case Studies on Artificial Intelligence Use in Policing, Public Defense, and Corrections (September 16, 2026), linking the police report-writing and corrections classification case studies.
8. Arizona Court of Appeals, Division One. State v. Horcasitas, No. 1 CA-CR 25-0191 (September 30, 2026).
9. ABC News. Attorney-general demands explanation from Tasmanian parole board over Susan Neill-Fraser AI mistake (September 16, 2026). See also Australian Associated Press, published by The Canberra Times (September 15, 2026, updated September 16) and the Newcastle Herald (September 18, 2026).
10. Sharon Bernstein. Northern California courts grapple with AI-generated errors: 1 prosecutor removed, Sacramento Bee (April 1, 2026).
11. Acknowledgment. The phrase "human in the lead, not human in the loop" was shared with the author by the AI initiative at Howard University School of Law.
12. Council on Criminal Justice Task Force on Artificial Intelligence. Principles for the Use of AI in Criminal Justice (October 2025); Assessing AI for Criminal Justice: A User Decision Framework (March 2026).
13. David Pitts and KiDeuk Kim. Responsible AI Adaptation in Corrections: Balancing Innovation, Oversight, and Human Dignity, Urban Institute (publisher page dated February 4, 2026).
14. Vera Institute of Justice. Accountability Principles for Adopting Artificial Intelligence Within the Criminal Justice and Immigration Systems (2026).
15. Nevada Legislature. NRS 176.0136, subpoenas; NRS 176.0137, agency cooperation; NRS 176.01385, confidentiality.
16. California Legislature. Penal Code §13663, added by SB 524, Stats. 2025, ch. 587, effective January 1, 2026. See also the San José Police Department's April 27, 2026 order reproducing the statutory requirements.
17. U.S. Department of Education. Prison Education Programs Questions and Answers, PEP-Q31, FERPA applicability and disclosure exceptions.
18. Nevada Legislature. SB 199 (2025): introduced bill; first reprint with amendments adopted April 21, 2025, section 18; NELIS history and overview.
19. Nevada Current. Lawmaker seeks to establish guardrails, 'some accountability' around artificial intelligence (April 3, 2025).
20. Legislative Counsel Bureau, Research Division. Reports of Certain 2025–2026 Interim Committees to the Legislative Commission (September 24, 2026), Bulletin 27-6, Joint Interim Standing Committee on Government Affairs, recommendation 9, printed page 40 (PDF page 46), BDR –504.
21. Nevada Legislature. NELIS 84th Session Bill Draft Request List, checked October 28, 2026: BDR –504.
22. Nevada Legislature. NRS 176.0134.