

Answering The ECCP's AI Questions About Your Speak-Up Channel.

The September 2024 ECCP asks prosecutors to probe both how a company governs its use of AI and whether its reporting mechanism actually works. A speak-up channel that uses AI sits at the intersection. This paper maps the questions to the evidence.

Published	September 2026
Edition	1.0
Source	Evaluation of Corporate Compliance Programs, updated September 2024
Audience	Chief compliance officers, GCs, internal audit
Author	XPOSER.AI

This is not legal advice. It is a practitioner's reading of a public Department of Justice document, written by a vendor whose product is one of the things it discusses. Where our own platform is relevant we say so explicitly rather than implying a neutral voice. Take the framework, verify the quotes against the source, and have your own counsel apply it.

CONTENTS

01 Why 2024 changed the question — three revisions, one intersection

02 What the ECCP actually asks about AI — the exact wording

03 What the ECCP actually asks about reporting channels

04 The intersection: an AI-assisted speak-up channel

05 The evidence file — what to be able to produce

06 Six failure modes prosecutors are equipped to find

07 A readiness self-assessment

08 How XPOSER.AI maps to these questions

09 References

Section 01

Why 2024 Changed The Question

On 23 September 2024 the Criminal Division published its fifth revision of the Evaluation of Corporate Compliance Programs. Three of its changes matter here, and they converge on a single system.

The revisions addressed, in the Department’s framing: risks associated with emerging technologies including artificial intelligence; encouragement of internal reporting and protection of whistleblowers; and whether compliance functions have adequate resources and access to data.

Read individually, each is a familiar compliance topic. Read together, they describe an uncomfortable convergence for any organisation that has recently modernised its hotline. The speak-up channel is simultaneously:

- the mechanism whose *effectiveness* the ECCP tests directly;
- increasingly, a deployment of *AI* whose governance the ECCP now tests;
- and the primary *data source* the compliance function must be able to reach.

A company can answer the AI questions well for its trading systems and still answer them badly for the system that triages allegations of misconduct — which is precisely the system a prosecutor examining a failed compliance programme will be looking at.

The practical consequence

If your hotline vendor introduced AI classification, summarisation, translation or routing, your programme acquired an AI governance obligation on the day that feature was switched on. The ECCP does not distinguish between AI you built and AI you bought.

Section 02

What The ECCP Actually Asks About AI

The technology questions sit in Part I, under the risk assessment discussion. The wording matters, so it is reproduced rather than paraphrased.

“How does the company assess the potential impact of new technologies, such as artificial intelligence (AI), on its ability to comply with criminal laws?”

ECCP, September 2024 — risk management

“Is management of risks related to use of AI and other new technologies integrated into broader enterprise risk management (ERM) strategies?”

ECCP, September 2024

“What is the company’s approach to governance regarding the use of new technologies such as AI in its commercial business and in its compliance program?”

ECCP, September 2024

Note the last five words. The Department distinguishes AI used to run the business from AI used to run the compliance programme, and asks about both. An AI-assisted intake and triage system is squarely the second.

The remaining technology questions are the operational ones:

- “Do controls exist to ensure that the technology is used only for its intended purposes?”
- “What baseline of human decision-making is used to assess AI?”
- “How is accountability over use of AI monitored and enforced?”
- “How does the company train its employees on the use of emerging technologies such as AI?”

“What baseline of human decision-making is used to assess AI?”

This is the hardest question in the set and the one most companies cannot answer. It presumes you know how your humans performed at the same task, because otherwise “the AI agrees with our reviewers 94% of the time” is a number without a denominator. A programme that cannot state its human baseline cannot demonstrate that introducing AI improved anything — or that it did not quietly degrade something.

Section 03

What The ECCP Asks About Reporting Channels

Part II, section D — *Confidential Reporting Structure and Investigation Process* — opens by describing what a well-designed programme has:

“... an efficient and trusted mechanism by which employees can anonymously or confidentially report allegations of a breach of the company’s code of conduct, company policies, or suspected or actual misconduct.”

ECCP, September 2024 — Part II.D

It then divides into three sets of questions.

Effectiveness of the reporting mechanism

- “Does the company have an anonymous reporting mechanism and, if not, why not?”
- “How is the reporting mechanism publicized to the company’s employees and other third parties? Has it been used?”
- “Does the company test whether employees are aware of the hotline and feel comfortable using it?”
- “Does the company encourage and incentivize reporting of potential misconduct or violations? Conversely, does the company use practices that tend to chill such reporting?”
- “How does the company assess employees’ willingness to report misconduct?”
- “How has the company assessed the seriousness of the allegations it received?”
- “Has the compliance function had full access to reporting and investigative information?”

Commitment to whistleblower protection and anti-retaliation

- “Does the company have an anti-retaliation policy?”
- “Does the company train employees on both internal anti-retaliation policies and external anti-retaliation and whistleblower protection laws?”
- “To the extent that the company disciplines employees involved in misconduct, are employees who reported internally treated differently than others involved in misconduct who did not?”

- “Does the company train employees on internal reporting systems as well as external whistleblower programs and regulatory regimes?”

Properly scoped investigations by qualified personnel

- “How does the company determine which complaints or red flags merit further investigation?”
- “How does the company ensure that investigations are properly scoped?”
- “What steps does the company take to ensure investigations are independent, objective, appropriately conducted, and documented?”

The section rests on U.S.S.G. § 8B2.1(b)(5)(C), which contemplates a publicised system, which may allow for anonymity or confidentiality, through which employees may report or seek guidance without fear of retaliation.

Section 04

The Intersection

Put the two sets side by side and the overlap is not incidental. Each ECCP technology question has a specific meaning when the technology in question is the thing triaging misconduct allegations.

READING THE AI QUESTIONS AGAINST A SPEAK-UP CHANNEL

ECCP question	What it means for an AI-assisted hotline
Impact of AI on ability to comply with criminal laws	If misclassification buries a bribery allegation in a low-priority queue, the technology has directly degraded the programme’s core function. This is a risk assessment item, not an IT one.
Integration into enterprise risk management	The hotline’s AI belongs in the ERM register alongside other model risk — with an owner, a review cadence and a documented failure mode, not in a vendor’s release notes.
Governance of AI in the compliance program	Who approved the model? Who reviews its outputs? What happens when it is wrong? Can the compliance officer explain, without the vendor present, how a severity score is produced?
Controls ensuring intended purpose only	The intake model triages reports. It must not also be scoring reporters, inferring identity from speech, or feeding analytics that let a manager see who is complaining.
Baseline of human decision-making	You need a measured human baseline for triage accuracy and speed, from before or alongside the AI, or you cannot evidence that the AI helps. See section 7.
Accountability monitored and enforced	A named human owns every case decision, with an audit trail showing who accepted or overrode each recommendation and when.
Training employees on emerging technologies	Case handlers need training on what the model can and cannot do — including its failure modes. Reporters need to know an AI processes their report and a human decides.

The reframe worth taking away

AI in a speak-up channel is not primarily an efficiency story to tell the board. It is a control whose design, testing and human oversight a prosecutor may one day examine. Build the evidence while the programme is healthy; it cannot be assembled retrospectively.

Section 05

The Evidence File

The ECCP is a set of questions, and questions are answered with artefacts. These are the documents a programme running an AI-assisted channel should be able to put on the table.

ARTEFACTS TO MAINTAIN

Artefact	What it evidences	Owner
AI risk assessment for the intake channel	The first three ECCP technology questions	Compliance + risk
Entry in the ERM register	Integration into broader risk management	Risk
Model governance note: purpose, inputs, outputs, limits	Governance and intended-purpose controls	Compliance
Human baseline measurement, dated	“What baseline of human decision-making...”	Compliance
Periodic AI/human agreement reporting	Monitoring and accountability over time	Compliance
Case audit trail export	Independent, documented investigations	Compliance
Override log with reasons	That humans genuinely decide	Compliance
Awareness and trust survey results	“Does the company test whether employees... feel comfortable using it?”	HR + compliance
Anti-retaliation policy and training records	The whistleblower protection questions	HR + legal
Vendor DPA, subprocessor list, penetration test summary	Data access, security and third-party risk	Procurement + security

Section 06

Six Failure Modes

Each of these is visible in documents a company already produces, which is why they are worth naming.

1 · The unexamined upgrade

A vendor enables AI classification in a routine release. Nobody performs a risk assessment because nobody procured anything. The programme now runs a model no one governs.

2 · The absent baseline

Agreement rates are reported with no measured human performance to compare against. The number sounds reassuring and evidences nothing.

3 · Automation deference

Handlers accept the model’s severity in almost every case. Formally a human decides; in substance the model does. An override rate at or near zero is a warning sign, not a success metric.

4 · The chilling analytics feature

A dashboard shows reporting volume by department and shift. In a small team that is close to identifying reporters — and it is exactly the kind of “practice that tends to chill such reporting” the ECCP asks about.

5 · Access asymmetry

The vendor holds analytics the compliance function cannot reach without raising a support ticket. The ECCP asks whether compliance has *full* access to reporting and investigative information.

6 · The untested channel

The hotline is publicised, unused, and assumed healthy. Low volume is treated as an absence of misconduct rather than as an absence of trust, and no survey has ever asked employees which it is.

Section 07

A Readiness Self-Assessment

Fifteen questions. Each should be answerable in one sentence with a document behind it. Answering “we would have to ask the vendor” is itself a finding.

#	Question	Evidence
1	Does an AI risk assessment exist for the intake and triage channel?	Assessment
2	Is that risk in the ERM register with a named owner?	Register entry
3	Can compliance explain how a severity score is produced, unaided?	Model note
4	Is there a documented human baseline for triage accuracy and speed?	Baseline study
5	Is AI/human agreement measured on a fixed cadence?	Calibration report
6	What is the override rate, and is it plausibly non-zero?	Override log
7	Does every case decision carry a named human owner?	Audit trail
8	Are handlers trained on the model’s limits, not just its interface?	Training record
9	Are reporters told an AI processes reports and a human decides?	Channel notice
10	Is anonymous reporting available — and if not, is the reason written down?	Policy
11	Has employee awareness and comfort been surveyed in the last 12 months?	Survey
12	Could any analytics view identify a reporter in a small team?	Access review
13	Does compliance have direct access to all reporting data?	Access matrix
14	Is model training on your report data contractually restricted?	DPA clause
15	Are internal reporters treated differently at discipline than non-reporters?	Discipline record

Section 08

How XPOSER.AI Maps To These Questions

We are the vendor, so treat this section as a set of claims to test rather than as analysis.

ECCP concern	Our position
Baseline of human decision-making	Model and human classifications are compared and reported on a fixed cadence. The methodology, including how the human baseline is established, is published separately in <i>The 60-Second Triage Benchmark</i> so it can be criticised.
Governance and explainability	Every classification is delivered with the policy citations and reasoning behind it, so a handler — and later an examiner — can see why a report was scored as it was.
Accountability monitored and enforced	AI recommends; a human decides. Every view, assignment, override and decision is written to an append-only audit trail that the customer can export.
Intended purpose only	No speaker templates, no voice matching, no reporter scoring. Report content is not used to train models serving other customers. These are DPA terms, not marketing lines.
Anonymous reporting mechanism	Anonymous intake with a zero-knowledge two-way follow-up channel, so anonymity does not cost the investigation its follow-up questions. Architecture detailed in our <i>Anonymity Whitepaper</i> .
Compliance access to data	Case data, audit trail and reporting are available directly to the customer's compliance function without vendor mediation.
Practices that chill reporting	Analytics are designed to avoid cohort sizes small enough to identify a reporter. If you can construct a view that does, tell us — we treat it as a defect.

What we cannot do for you

Anti-retaliation policy, discipline consistency, training, publicising the channel and surveying whether employees trust it are programme obligations no vendor can discharge. A tool can make a good programme demonstrable; it cannot make an absent one exist.

Section 09

References

1. **Evaluation of Corporate Compliance Programs**, U.S. Department of Justice, Criminal Division, updated September 2024. Quotations in sections 2 and 3 are taken from this document. justice.gov/criminal/criminal-fraud/compliance
2. **U.S. Sentencing Guidelines** § 8B2.1(b)(5)(C), cited within the ECCP in support of the confidential reporting discussion.
3. **OMB Memorandum M-24-10**, which supplies the definition of “artificial intelligence” adopted by the ECCP at footnote 4.
4. **Principles of Federal Prosecution of Business Organizations**, Justice Manual 9-28.000.
5. **XPOSER.AI, Anonymity Whitepaper** and **The 60-Second Triage Benchmark: Methodology**, xposer.ai/resources

XPOSER.AI · Autonomous, voice-first AI whistleblowing. From incident to audit-ready insight in under 60 seconds — 24/7, in 100+ languages.

Questions or corrections: info@xposer.ai

© 2026 XPOSER.AI. Provided for information only and not legal advice. Quoted material is from a public U.S. Government document and is reproduced for commentary; verify against the source before relying on it.