



ACF IN ACTION

acfstandard.com

ACF en Action applique l'Agentic Commerce Framework® à des entreprises réelles, à partir de sources publiques. Chaque cas montre qui reste responsable quand l'IA décide.

ACF-CS-021 · Case Study #021 · Édition 2026

NYC Local Law 144

How do you delegate the screening of people to an agent while making bias visible and controllable?

CARTE D'IDENTITÉ

SECTOR Regulation / HR	COUNTRY USA (NYC)	TECHNOLOGY Automated screening (AEDT)	TYPE Framed	AUTONOMY framed
CONFIDENCE High				

OUTILS ACF ILLUSTRÉS DANS CE CAS

ACF-01	ACF-03	ACF-05	ACF-08	ACF-00	ACF-02
--------	--------	--------	--------	--------	--------

mobilisé · à prévoir · 6 outils du Toolkit (4 mobilisés, 2 à prévoir, sur 17).

1. Executive summary

Since July 5, 2023, the City of New York has enforced Local Law 144: any organization that uses an Automated Employment Decision Tool (AEDT) to hire or promote must first have it audited for bias. The delegated decision is that of a screening agent: to recommend setting aside or retaining candidates. The governance approach is distinctive: instead of leaving the bias invisible, the law requires it to be measured and published.

Read through the lens of ACF, this case is a reference case on the bias of a delegated decision. Where an opaque deployment lets the screening operate in the shadows, Local Law 144 imposes an independent, annual, and public audit. The value of ACF here is prospective and confirming: the law materializes, in law, what the framework calls fairness supervision. Bias is not assumed absent, it is made visible so it can be corrected.

2. Context

Like most large organizations that recruit at scale, New York employers rely on automated tools to screen applications and rank profiles. The intent is legitimate: to handle a high volume without dedicating a human recruiter to every file. The question the case raises is just as legitimate: when this tool recommends setting aside or retaining people, how do you ensure it does not reproduce a protected bias without anyone seeing it?

Local Law 144 answers with a procedural constraint: no deployment without an independent, published fairness measure.

3. Public sources used

- City of New York (DCWP), "Automated Employment Decision Tools (AEDT)" page.
- Deloitte, "NYC Local Law 144-21 and Algorithmic Bias".
- National Law Review, "NYC's Local Law 144 and the Final Regulations".

4. Reconstructed AI architecture (from public sources)

Mechanism imposed by Local Law 144 (in force since July 5, 2023).

- Nature: a legal framework governing the use of automated employment decision tools (AEDT) for hiring and promotion.
- Decision concerned: to recommend setting aside or retaining a candidate, that is, to screen people.
- Control function: to make bias measurable and enforceable, rather than to prohibit it or leave it in the shadows.

The documented obligations:

- an independent and annual bias audit, carried out by a third party, on historical or test data;
- a measure of disparate impact by protected characteristics (sex, race or ethnicity, and their intersections);
- publication of an audit report setting out the method and the impact results;
- notice to candidates that an automated tool is being used.

What the mechanism makes possible:

- to see the bias, through an impact measure and not an intuition;
- to correct it, from a quantified and dated starting point;
- to answer for it, through a public report and an independent third party.

Non-compliance exposes the organization to penalties (up to 1,500 dollars per violation per day).

5. ACF mapping

 **Tools used:** ACF-05 (fairness supervision) · ACF-03 (agentic constitution) · ACF-08 (register).

In ACF language, the mechanism reads as follows:

ACF element	In Local Law 144
Named human (accountable)	The employer that deploys the AEDT: it is they who answer for the screening and its compliance with the audit obligation.
Agent	The automated employment decision tool (AEDT) that recommends setting aside or retaining candidates.
Platform	The underlying screening system, provided or integrated, on which the employer relies.
Delegated decision	The screening of people: to rank, set aside, or retain applications.
Autonomy level	Bounded: the tool may screen, but only if it has been audited for bias and the result is published.
Non-delegable zones	Deploying a screening of people without an up-to-date fairness measure: the law in effect prohibits it.
Rule / safeguard	Independent and annual audit, measure of disparate impact, publication of the report, notice to candidates.
Control third party	The independent auditor who carries out the measure, and the regulator (DCWP) who penalizes non-compliance.
Traceability	The public audit report: method, impact results, dates. It constitutes the enforceable trace of measured fairness.

A schema in ACF graphic language accompanies this case (separate file).

6. Analysis

 **Tools used:** ACF-05 (supervision) · ACF-03 (non-deployment rule) · ACF-08 (register). **To link:** ACF-01 (decision map) · ACF-02 (criticality).

What the mechanism does well (measurement rather than assumption):

- **Fairness becomes a controllable fact.** The law does not ask you to trust the tool, it requires measuring its fairness and publishing the result. Bias stops being a blind spot and becomes a quantified and dated piece of data.
- **Independence and publication commit.** An audit carried out by a third party and made public cannot be reduced to a self-declaration of non-bias. This is what makes the measure credible and enforceable.
- **Annual repetition embeds fairness over time.** The audit is not a one-off control: its renewal each year installs continuous monitoring, not a single snapshot.

What ACF adds and formalizes:

- **A constitutional rule.** Do not deploy a screening tool without an up-to-date fairness audit: this is the translation, into ACF language, of the legal obligation into a permanent and signed rule.
- **A usable register.** Keep the audit reports, the method, the impact results, and the dates, so as to answer for them and track the evolution of bias over time.
- **A loop to close.** Measuring is the condition, but correcting and then republishing closes the loop. The annual audit materializes this continuity: a decision about people delegated without a fairness measure remains a governance blind spot.

7. ACF toolkit used

This grid indicates the tools that the ACF analysis uses (✓) to read this case, and those to link (→
SOON). It does not mean the organization uses ACF: it shows which instruments the framework would bring.

Tool	Status	Role in this case
ACF-05 · Supervision	✓	Describe the fairness audit mechanism (indicator, third party, publication)
ACF-03 · Agentic constitution	✓	Set the rule "no deployment without an up-to-date fairness audit"
ACF-08 · Decisions register	✓	Keep reports, method, impact results, and dates
ACF-01 · Decision map	→ SOON	Locate the screening agent, the accountable party, and the delegated decision
ACF-02 · Criticality matrix	→ SOON	Classify the screening of people as a high-criticality decision
ACF-00 · Sovereignty Score	→ SOON	Assess the maturity of the measurement mechanism

8. Governance questions (what an AI committee should ask itself)

1. Why does an independent and public audit strengthen trust in a screening agent, more than a self-declaration of non-bias?
2. Is measuring fairness enough, or must you also correct and republish to close the loop?
3. Can the screening of people be delegated to an agent without any fairness measure?
4. Are our candidates informed that an automated tool takes part in their evaluation?
5. Do we hold a dated register of fairness measures, reconstructable and enforceable?

9. General lessons

A screening agent does not fail only when it is biased: it fails when it is allowed to screen without ever measuring that bias. Local Law 144 does not constrain the tool, it requires that its fairness be measured, published, and controlled by a third party. What makes the delegation defensible is not the promise of a neutral tool, but the proof, independent and dated, of its fairness.

The transferable lesson: delegating the screening of people is acceptable only if the bias is measured, independently and publicly. Measuring makes the delegation defensible; assuming the absence of bias never does.

10. What ACF would do (reference design)

Prescriptive and non-critical formulation: here is how an organization would design such a mechanism directly with the framework, whatever its sector.

If an organization wished to deploy a screening of people directly with ACF, it could in particular:

- **map** the screening decision on the autonomy scale (ACF-01) and classify it as high criticality, since it bears on people;
- **declare** in a signed agentic constitution (ACF-03) the rule: never deploy a screening tool without an independent and up-to-date fairness audit;
- **organize** a fairness supervision (ACF-05): disparate-impact indicator, measurement by a third party, annual publication;
- **document** a decisions register (ACF-08) linking each deployment to its audit report, its method, its results, and its dates;
- **close the loop** by correcting and then republishing, rather than treating the audit as a one-off formality.

None of this is a reproach: it is the way ACF would make **explicit, signable, and auditable** what the law already imposes, by embedding the fairness measure in continuous governance.

11. Similar cases (transferability)

This case is specific neither to New York nor to recruitment. It is transferable to any organization that delegates to an agent a screening that bears on people:

-  human resources and recruitment
-  credit and bank scoring
-  insurance (pricing, risk selection)
-  admission in education
-  allocation of benefits or social welfare
-  moderation and ranking of internal applications

Everywhere, the same pattern: an agent that ranks people, a possible bias to be measured rather than assumed, and an independent trace to be published. This is what makes this document a reference case: it illustrates a concept, not just a local law.

12. Confidence of the assessment

High. The case rests on a law in force (NYC Local Law 144) and its implementing regulations, documented by the official source (DCWP) and by analyses from leading firms. The facts, obligations, and dates are therefore established and verifiable by the reader. The ACF analysis draws on these public facts; it bears on the governance of bias measurement, not on the internal technical details of the tools concerned.

Teacher guide

Pedagogical objective. Convey that a screening of people delegated to an agent is defensible only if the bias is measured, independently and publicly, rather than assumed absent.

Session outline (90 min).

1. (15 min) Reading of the case, without the Analysis part.
2. (20 min) In groups: describe the fairness audit mechanism in ACF language (indicator, third party, publication, card ACF-05).
3. (20 min) Workshop: write the rule "no deployment without an up-to-date fairness audit" and the content of the register (cards ACF-03 and ACF-08).
4. (20 min) Debate: "Can a screening of people be delegated without a fairness measure?" (expected answer: no, it is a governance blind spot).
5. (15 min) Synthesis by the teacher: measuring, independently and publicly, makes the delegation defensible.

Possible exam questions. See the "Governance questions" section.

Reminder: grading students' work is a pedagogical act of the teacher. This guide provides the reference analysis, not automatic grading.

Associated lab (ACF-LAB-021)

From this case, the student produces:

- the fairness supervision mechanism (card ACF-05): impact indicator, third-party auditor, publication method;
- the constitutional rule (card ACF-03): do not deploy a screening tool without an up-to-date fairness audit;
- the schema of the register (card ACF-08): reports, method, impact results, dates to keep;
- a one-page comparison between a screening with invisible bias and a screening with measured and published bias.

Appendix - Completed ACF cards

The official Toolkit cards used in this case, filled in with its public data. For illustration: an organization of the same profile can reuse and adapt them.



OBJECTIVE

Map the decisions delegated to the screening tool (AEDT) and set, for each, the autonomy level based on what it commits on people, not on the volume processed. Filled from public sources.

1 AGENTIC MATURITY SCALE

0

Classic automation

Fixed rules, no autonomy.

1

Assisted agents

The agent proposes, the human decides.

2

Governed agents

The agent acts within a frame, under supervision.

TARGET ~80%

3

Supervised autonomy

The agent decides alone; control after the fact.

2 YOUR KEY DECISIONS (filled)

A

Pre-filter on an explicit criterion

STAKE

Check an objective condition (required degree, availability)

IMPACT

Reversible, reviewed by a recruiter

DEADLINE

Continuous

AUTONOMY

Autonomous

B

Rank applications

STAKE

Sort profiles for a human recruiter

IMPACT

Orients without rejecting on its own

DEADLINE

Continuous

AUTONOMY

Supervised autonomous

C

Reject or retain people without a fairness audit

STAKE

Screen people on an unmeasured model

IMPACT

Commits fairness and the employer's accountability

DEADLINE

-

AUTONOMY

NON-DELEGABLE

Reading: it is this setting, made **before** deployment, that distinguishes a tooled screening from a screening of people delegated without a fairness measure, which Local Law 144 effectively prohibits.



OBJECTIVE

Define the agent's constitution: identity, standing principles and forbidden zones. Card filled with public case data.

1 AGENT IDENTITY & MISSION

AGENT NAME

Automated employment decision tool (AEDT)

ACCOUNTABLE (DDAO)

The employer that deploys the tool

PRIMARY MISSION

Screen and rank applications at scale

SCOPE

Screening of people; subject to an independent, published bias audit

2 CORE PRINCIPLES

- 1 Never deploy a screening of people without an up-to-date fairness audit
- 2 Measure disparate impact across protected characteristics
- 3 Publish the audit report: method and results
- 4 Inform candidates that an automated tool is involved

EXAMPLE PRINCIPLES

- Respect the legal and sector framework
- Protect data and confidentiality
- Stay transparent, every decision traceable
- When in doubt, escalate rather than decide

3 FORBIDDEN ZONES (ABSOLUTE)

- 1 Deploy or keep the tool without an annual independent bias audit
- 2 Treat the audit as a one-off formality, without correction or republication
- 3 Screen people on a model whose bias is not measured

EXAMPLE PROHIBITIONS

- X Commit spending above a threshold without validation
- X Share personal data with a third party
- X Change its own rules or scope on its own

Constitution filled from the text in force (NYC Local Law 144). Principle 1 translates into a standing, signed rule the legal duty not to deploy without a fairness measure.



OBJECTIVE

Organize fairness supervision: measure the bias, publish it, correct it over time. Filled with the case's publicly documented obligations.

1 SUPERVISION INDICATORS

Indicator	Value / target (public sources)
Bias audit	independent, annual, by a third party
Disparate impact measure	by sex, race/ethnicity and intersections
Report publication	method and results made public
Candidate notice	an automated tool is used
Penalty for breach	up to \$1,500 per violation and per day

2 SETUP

- 1 Disparate impact indicator measured on historical or test data
- 2 Measure performed by an independent third party, renewed each year
- 3 Publication of a report setting out method and results
- 4 Closed loop: correct then republish, not a one-off check

Measurement is the condition, not the end. Linking the fairness audit to the register (ACF-08) places the measured bias under continuous monitoring.



OBJECTIVE

Keep every fairness audit to be able to answer for it and track the bias over time. Filled with the fields relevant to the case.

1 FIELDS TO RECORD

Field to record	Example, Local Law 144 case
Timestamp	Date of the annual audit
Tool concerned	AEDT screening applications
Audit method	Disparate impact measure on test data
Impact results	Ratios by sex, race/ethnicity and intersections
Publication	Public audit report, dated

2 USE OF THE REGISTER

- 1 Track the bias evolution from one audit to the next
- 2 Correct from a quantified, dated baseline
- 3 Establish who answers for the screening: the employer
- 4 Provide enforceable evidence of the measured fairness

Without a dated register of fairness measures, the audit stays a point-in-time snapshot. The register turns it into enforceable evidence and continuous monitoring.

Disclaimer

This study is carried out **exclusively from public information** (a law in force and accessible legal analyses). It constitutes **neither an audit, nor a certification, nor legal advice**. It applies the ACF framework for pedagogical and illustrative purposes. No internal, confidential, or non-public data was used. The trademarks and institutions cited belong to their respective owners.

Sources

- City of New York (DCWP), *Automated Employment Decision Tools (AEDT)* - <https://www.nyc.gov/site/dca/about/automated-employment-decision-tools.page>
- Deloitte, *NYC Local Law 144-21 and Algorithmic Bias*
- National Law Review, *NYC's Local Law 144 and the Final Regulations*

ACF Case Study ACF-CS-021 · v1 · working document, to be reviewed before any public distribution.