



AI Compliance in Business and Technology

What goes wrong, what the law requires, and a 90-day plan to get ahead of it

Webinar | 2026

Agenda

Incidents	07
Landscape	17
90-Day Plan	27
What You Get	42



About the Speaker

Sergii Shelpuk

- Worked in AI & Data Science since 2013
- Led AI & Data Science departments for 4 companies, including a billion-dollar one
- Helped raise \$52.5M in venture founding as a management team member
- Co-founded an AI product for biopharma R&D
- Graduated from LSE and Stanford



Sergii Shelpuk

Edvantis AI Consulting Expert,
Global Strategy and
Implementation

About Edvantis

Edvantis is a global software engineering company helping businesses stay competitive in the digital world.

We modernize infrastructure, automate processes, and deliver AI-driven solutions to improve efficiency, agility, and growth.



20+ years of experience



400+ employees



450+ delivered projects



40+ active clients



ISO/IEC 27001:2013
Certified



ISO 9001:2015
Certified



2025 | Clutch Global
Awards Winner



Software Development
Services Grid Leader | G2



Why Compliance Matters

Compliance is seen as a drag.

In reality it protects your business
and your customers.

\$25M

lost to a single deepfake call –
details ahead

Question to the Audience

Does your company have an AI use policy?

☐

Yes, a formal one

☐

No policy

☐

Informal rules only

☐


What's AI?



01/ AI Risks

What AI mistakes cost other companies

Replit / SaaStr: Deleted Production Database



Source: [Fortune](#)

1,206

executive records wiped in
one command

The setup

Public 'vibe coding' trial with explicit instruction not to change code.

The cover-up

AI Agent claimed rollback was impossible. It wasn't.

The response

Replit's CEO called it "unacceptable". Shipped dev/prod database separation.

ⓘ An agent with production credentials and no approval gate is one bad inference away from an outage.

EEOC v. iTutorGroup: First Regulator Settlement

Press Release

09-11-2023

iTutorGroup to Pay \$365,000 to Settle EEOC Discriminatory Hiring Suit

Source: [Equal Employment Opportunity Commission](#)

200+

qualified applicants auto-rejected

What the software did

Auto-rejected women 55+, men 60+

How it was caught

Same CV, younger birth date → interview.

The outcome

\$365K consent decree + five years of EEOC monitoring.

ⓘ The company never knew it was rejecting qualified people. It surfaced only because one person tested it.

Mobley v. Workday: The Vendor on the Hook

Workday must face California lawsuit over AI bias in job screening tools

By Daniel Wiessner

June 22, 2026 11:12 PM GMT+3 · Updated June 23, 2026



Source: [Reuters](#)

100+

job rejections behind one
discrimination lawsuit

The claim

Rejections via Workday's tools – race, age, disability.

The scale

Nationwide ADEA collective certified, applicants 40+.

The spillover

Workday forced to disclose which customers enabled HiredScore AI features.

ⓘ Still ongoing, and aimed at the vendor – with the vendor's customers pulled into litigation they never joined.

Moffatt v. Air Canada: Chatbot Speaks for You

Air Canada ordered to pay customer who was misled by airline's chatbot

Company claimed its chatbot 'was responsible for its own actions' when giving wrong information about bereavement fare

Source: [Guardian](#)

CA\$812

total damages and a precedent every chatbot case now cites

What happened

Chatbot promised a discount policy that didn't exist.

The defence that failed

"Not liable for our chatbot" – tribunal disagreed.

The award

CA\$650.88 in damages plus interest and fees.

ⓘ The award was small, but the principle matters – whatever your chatbot promises, your company promised.

Deloitte Australia: Fabrications in a Paid Report

Deloitte was caught using AI in \$290,000 report to help the Australian government crack down on welfare after a researcher flagged hallucinations

By Nino Paoli
Former News Fellow

Follow

October 7, 2025, 5:10 PM ET



Source: [Fortune](#)

A\$97K

refunded after finding AI-generated errors in the report

The deliverable

A\$440K government report – fabricated quote, nonexistent sources.

The disclosure

Revised version admitted Azure OpenAI use.

The refund

~A\$97K – a quarter of the contract.

ⓘ The refund wasn't the real cost. The reputational damage was – every client now has a reason to doubt the deliverables.

Arup: \$25M Deepfake Video Call

EMERGING TECHNOLOGIES

‘This happens more frequently than people realize’: Arup chief on the lessons learned from a \$25m deepfake crime

Feb 4, 2025

Source: [World Economic Forum](#)

\$25M

lost to a single
deepfake call

The attack

15 transfers totalling \$25M – every other caller, incl. the CFO, was a deepfake.

No breach

No systems compromised. Pure social engineering.

The outcome

Money never recovered.

ⓘ “I saw and heard them” no longer proves anything. The control is payment auth over a separate channel.

More of the Same Pattern (1/2)

Amazon recruiting engine (2018)

Trained on ten years of male-skewed CVs, learned to penalize the word “women’s”.

Chevrolet of Watsonville (Dec 2023)

Dealer chatbot persuaded into agreeing to sell a Tahoe for \$1 with a “legally binding” confirmation.

NYC MyCity chatbot (2024)

Advised business owners in ways that would have broken city law.

Cursor support bot (Apr 2025)

Invented a device-limit policy that didn’t exist – triggering public cancellations.

More of the Same Pattern (2/2)

Samsung (2023)

Engineers pasted proprietary source code into ChatGPT; an internal ban followed.

Mata v. Avianca (2023)


Lawyers sanctioned for filing ChatGPT-fabricated case citations.

Zillow Offers (2021)

Algorithmic home-pricing wound down with a ~\$304M writedown and ~25% of staff cut.

Bottom Line

Regulations are designed to avoid these and similar incidents.



02/

The Regulatory Landscape

Two fundamentally different regimes: EU and US

EU vs US At a Glance



EU


- One horizontal law
- Rules apply before the system ships
- Risk-tiered obligations
- GDPR + AI Act as a coherent pair




US

- No horizontal AI law
- Liability comes after the harm – general law, sector regulators, courts
- A state patchwork on top – and an active federal effort to flatten it
- No federal privacy law either: sectoral + state privacy statutes

EU AI Act – Who It Applies To




GDPR – Extraterritorial Scope



Any one of these = GDPR applies to you

US – Three Layers

Some regulation for model developers and providers – very little for AI users.



Layer 1 – Existing Law, Applied to AI

No AI statute needed – this is where enforcement actually happens.

False AI claims - FTC Act §5

AI-washing, overstated capabilities

Sector rules catch AI - HIPAA / FDA / SEC / banking

Vertical regimes that bind AI without naming it

Hiring discrimination - Title VII / ADEA / ADA

AI screening tools already getting sued

Chatbot liability - State UDAP + contract / tort

Negligent misrepresentation, breach

Credit denials - ECOA / FCRA


Model says no – you must explain why



EEOC withdrew some AI guidance in January 2025 – Title VII, ADA and ADEA remain fully in effect.

Layer 2 – State AI Statutes

One country, many rulebooks. Texas, California, Illinois, NYC, Utah – each writing its own.




Layer 3 – Federal Executive Policy

Not adding rules – dismantling Layer 2 through courts and funding.

- 
- Jan 2025** ● Biden's AI executive order revoked. New order sets a deregulatory course.
 - Jul 2025** ● Senate strips a 10-year moratorium on state AI laws from the budget bill, 99–1. Congress declines to revive it.
 - Dec 2025** ● DOJ task force to challenge state AI laws in court. \$42B in broadband funding tied to states repealing AI rules. Exempt: child safety, data centers, state procurement.
 - Mar 2026** ● White House asks Congress to preempt state AI laws imposing "undue burdens."
 - Jul 2026** ● FTC proposes a policy statement on "suppression of accuracy" in AI systems.

The Colorado Story

The only state law that genuinely copied the EU AI Act's approach – risk tiers, impact assessments, anti-discrimination duties. Then—



Question to the Audience

Do you know how many AI systems run in your company right now?

☐

Yes, exactly

☐

Roughly

☐

No idea

☐

Someone probably knows




03/

Compliance Best Practices

A month-by-month 90-day AI governance plan

AI Governance in **Your Organization**

It makes you safe, builds customer trust,
and makes you compliant – **all at once.**



Month 1 – the AI Census

Find every system where AI shapes a decision.

1

person

part-time, not auditing

4

weeks

every system with AI output

1

spreadsheet

systems + owners



The rule: You are not auditing anyone. Say that first, every time you ask a question.

Week 1 – Pull Records

Four sources. No interviews yet.

Finance	Invoices and card charges containing "AI", "GPT", "Copilot", "assistant".
IT	SSO logs: which apps people actually sign into
Engineering	API keys issued – and which model providers they point to
Procurement	contracts from the last 18 months where the vendor has since shipped AI features

System	Source	Found via
Claude Code	Engineering	IT + Finance
Copilot	Sales	Procurement
Midjourney	Design	Finance

this is what you get at the end




Expect the expense cards to surface tools IT doesn't know about. That's normal – log it, don't escalate it.

Week 2 – Ask Department Heads

Three questions, in this order. Don't use the word "AI" until they do.

1. Where does software shape a decision or a deliverable someone relies on?
2. What happens to that output – a person checks it, it goes straight to a customer, or the system acts on it?
3. What got noticeably faster or cheaper here in the last two years – and why?



Week 3 – One Row Per System

Eight columns. Resist adding more.

System	Owner	What it produces	Data it touches	Built or bought	Affects a person's outcome	Acts or only suggests	Where users are located
CV screening	K. Adams, HR	Candidate ranking	CVs, contacts	Bought	Yes, hiring	Suggests	US + EU
Support chatbot	...	Customer replies	...	...	...	Acts	...
Invoice extraction	...	...	...	...	...	...	...



Incomplete rows are fine. Missing systems are not.

Week 4 – Sort and Assign

RED

Affects someone's life outcome – employment, credit, housing, healthcare.
Or it acts on production without approval.

AMBER

Touches personal data, or its output reaches a customer.

GREEN

Internal, no personal data, a human checks everything.



Every Red system gets a named owner. A person, not a team.

What You Deliver at Day 30



Deliverables

1. The spreadsheet
2. A count you can state in a board meeting
3. A named owner per Red system



What to expect

- The count comes in three to five times higher than the estimate
- The highest-risk system is usually one nobody bought as AI – it arrived as a feature in a tool you already had

What the **Census Buys You**

1.

GDPR paperwork

Your data-processing register probably exists already. The census fills in the AI rows.

2.

AI Act obligations

Which of your systems are in scope – and whether you're a provider or a deployer.

3.

Banned-practices check

Some AI uses are outright banned. One day to check your list against theirs.

4.

Training needs

Who in the company needs AI literacy training – the session itself is a separate week.

Week 5 – Prohibited-Use Screen

Take the Red and Amber lists. Check every system against the AI Act's banned practices:



Social scoring



Manipulative or
exploitative techniques



Emotion recognition at
work or in education



Untargeted face-image
scraping



Biometric sorting by
protected traits



Predictive policing of
individuals



One day of work. Penalties reach 7% of global turnover. Document that you ran the check – the record is the deliverable.

Week 6 – Production Access

One question to engineering:

Which systems can change or delete production data, send money, or contact customers – without a human approving it?

That list should be **empty**. Gate everything on it.



This is the Replit control. It has no legal deadline – you do it because you don't want to explain a deleted database to your board


Weeks 7–8 – Train People

One 45-minute session with attendance recorded. Three rules to teach.

1. What never goes into an external tool?
2. What always gets checked before it leaves the building?
3. Who to tell when something looks wrong?



By day 60 you've closed both AI Act requirements that already apply – the banned-uses check and staff training.



Weeks 9–10 – Disclosure

Walk your own customer journey. Find what's hidden.



Talking to a machine

Chatbot, voice agent, auto-reply – the customer doesn't know it's not a person.



AI-generated content

Product descriptions, summaries, reports – nothing says it was generated.



Automated decision

Pricing, eligibility, ranking – a decision was made about them and they weren't told.



Before deploying AI on your own workforce, inform workers' reps. Companies forget – and it applies to exactly the HR tools they're buying most.

Weeks 11–12 – Contracts & Decisions

Two checks for every Red and Amber system.

1. The vendor contract

☐ Is there a data-processing agreement?

☐ Does it stop the vendor training on your data?

☐ Where is the data hosted?

2. The automated-decision flag

Does any system make a decision about a person – entirely on its own?

→ If yes, separate GDPR obligations apply.



By day 90 you've closed the AI Act transparency duty (binding since 2 August 2026), plus your GDPR transparency and contract position.

After Day 90 – Where the Work Gets Heavy

Months 4–12 stop being a sprint. Five workstreams ahead:



Impact assessments

Proof you assessed the risks before launch



Human oversight

A human who can stop the system



Testing

Bias, accuracy, drift



Logging

Why did it decide that



Incident playbook

Who to notify, and how fast

2 Dec 2027

High-risk obligations land



Roughly 15 months of work from a standing start. The deadline is closer than it looks.

edvantis

04/

What You Actually Get

Nine business wins. Compliance comes free.

What Doesn't Happen

Headline	Case	Takeaway
You don't lose your secrets	Samsung engineers pasted proprietary source code into ChatGPT. No recall, no containment.	Data governance is IP protection with a compliance label on it.
You don't wake up to a deleted database	Replit's agent wiped a live database during a code freeze – then claimed it was irreversible.	Least privilege and dev/prod separation aren't AI controls. You already claimed to have them.
You don't pay a fraudster	Arup lost \$25M on a call where every participant except the victim was a deepfake.	Payment auth over a separate channel: dull, cheap, now non-optional.
You don't ship errors under your logo	Deloitte refunded part of A\$440K on a government contract – fabricated citations, nonexistent sources.	Review gates protect the thing you actually sell: being trusted.
You stop losing good candidates	iTutorGroup rejected 200+ qualified applicants. It surfaced because one person re-applied with a different birthdate.	Bias testing is a recruiting metric before it's a legal control.

What Does Happen


Headline	Point	Takeaway
You win deals that require answers	Enterprise buyers ask what AI touches their data, whether it trains on it, who's accountable.	Companies that answer in a day close faster than those that need weeks and a lawyer.
You know what you're paying for	AI spend hides in card charges, duplicate tools, forgotten subscriptions.	The census finds 3–5× more AI than estimated. The inventory pays for itself.
You adopt faster, not slower	Companies without governance ban AI after an incident. It goes underground.	Companies that know what they run approve the next project in a week.
When it goes wrong, you survive it	Not if – when. What matters is what you can show afterwards.	Can you see what the system did, prove you had controls, tell customers first?

None of This Is About AI

It's the same controls you'd demand of any part of your business that can act on its own.

Do it because you don't want to lose the candidate, leak the roadmap, wire the money, or explain the deleted database.

Compliance is what you get for free.



Question to the Audience

What will you do first after this session?

☐

Run the AI census

☐

Start staff training

☐

Get someone to help us with it

☐

Nothing yet – need to digest





Got a Question? Let's Talk!

Book an AI consulting session with our Team

Email us at: engagement@edvantis.com

Or talk to your Edvantis Account Manager.



What You'll Get:

AI Census Kickstart

We help you run the 30-day inventory from this webinar

Compliance Gap Check

Where you stand against what's already binding

A Realistic Roadmap

What to close now, what can wait

edvantis

Thank you!