HUB

Risk & Insurance | Employee Benefits | Retirement & Private Wealth

WEBINAR

AI in Employee Benefits:

Managing Compliance in an Era of Automated HR

Moderator

Nathan Wolfe

Employee Benefits Strategist

HUB International

Trevor Howe

Senior Benefit Technology Specialist and HCM Consultant

HUB International

Cory Jorbin, Esq.

National Director of Compliance Consulting, Employee Benefits

HUB International

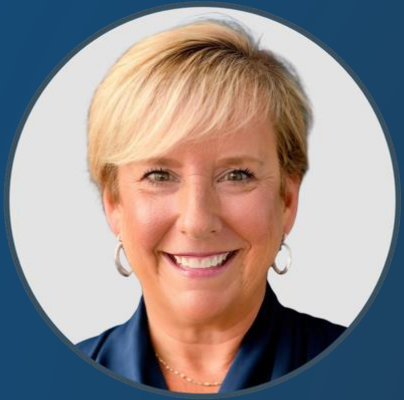

Michelle Jukoski

Executive Vice President, Technology Strategy & Workforce Solutions

HUB International

Agenda

1. AI is Already Part of Your HR Stack
2. Where AI Creates Compliance Exposure
3. ERISA Fiduciary Duty Doesn’t Transfer to Software
4. What an AI Governance Framework Looks Like
5. Using AI to Strengthen Your Team’s Capacity
6. Key Takeaways & Next Steps

Disclaimer

The information provided is solely for general information purposes and, in particular, is not intended as tax or legal advice. Any opinions expressed by the speakers are those of the speakers and not necessarily reflective of the opinion of HUB. We disclaim all liability and responsibility arising from any reliance placed on the information by you or anyone else who may be informed of any of its contents.

AI is Already Part of Your HR Stack

AI is Already in Your Stack: Four Types to Know

Your platforms added AI without a memo. Here's what's running.

Conversational AI

- Chat-style. Answers in plain language from approved sources.
- Already the front door of your OE platforms, HR portals, and ben admin tools

Predictive Analytics

- Flags turnover risk, utilization trends, workforce cost shifts — before they're problems
- Draws on historical data across multiple systems

Multimodal AI

- Works across text, voice, images, documents, and video in the same tool
- **Example:** Transcribes exit interview → summarizes themes → pushes to workforce analytics. No manual entry.

Agentic AI

- Takes a goal, executes the steps, follows up — without being re-prompted
- Proactive, not reactive. Acts on conditions.
- **Example:** Detects a tax law change → notifies you → tells you exactly how to update your system

The Agency Spectrum: Where Are You Right Now?

Gartner

60% of HR work tasks completed by AI agents by 2030. SHRM research (2026): only 41% of workers personally use AI for work today — and only 33% with formal organization backing.

Low Agency
Already in your stack

- OE chatbots, benefits FAQ responses, enrollment confirmations, performance review drafting
- **Watch:** Human is "in the loop" — but is anyone actually reviewing the output?

Medium Agency
Embedded in your HCM

- Benefits utilization analytics, sentiment analysis, ACA measurement tracking, OE drop-off detection
- **Watch:** Partial data → skewed results. AI misreads tone. Output influences decisions.

High Agency
Here or arriving fast

- Real-time payroll discrepancy detection, automated eligibility determination, multi-step onboarding automation
- Multiagent recruiting: source → schedule → interview → select. Agents hand off to agents.
- **Watch:** Governance must be built before this lands — not retrofitted after.

What the Research Shows:
SHRM 2026 AI Workplace Study

SHRM Navigating AI in the Workplace, 2026. Survey of 5,875 U.S. workers, March–April 2026.

Adoption: The Gap Between Hype and Reality

- **41%** of workers personally use AI for work
- **33%** use AI with formal organizational backing — the rest are on their own
- **47%** of AI-adopting organizations have an AI policy in place

The Policy Gap Creates Real Compliance Risk

- **30%** of AI-using workers admit to violating their org's AI policy
- **#1 violation:** Uploaded confidential company data to an unapproved tool (34% of violators)
- Only 61% of violators ever faced any consequence — the gap invites shadow AI

The Productivity Reality

- **68%** of AI users say it improved their efficiency — avg. 6 hours saved per week
- But 65% spend moderate to significant effort reviewing and correcting AI output
- Directors save 9 hrs/week. Individual contributors save 4 hrs — but pace expectations rise for all.

The Governance Warning

- **55%** of Directors & Above would follow AI's recommendation over their own ethical judgment
- Senior leaders — the group with the most decision-making authority — defer to AI most readily
- SHRM: deliberate practices needed to preserve independent judgment at the top

Source: Navigating AI in the Workplace: 2026, SHRM. Survey of 5,875 U.S.-based workers, March–April 2026.

Before AI vs. With AI: What the Shift Looks Like in Practice

Reactive → proactive. These are current or near-term vendor capabilities — not hypothetical.

	Payroll & Tax Compliance	Performance Management	Benefits Administration Eligibility & OE	Reporting & Analytics
Before AI	Receive a notice → research → correct retroactively → potentially incur penalties	Manager spends hours pulling peer feedback, comp comparisons, prior review notes	Manual eligibility audits. Enrollment errors found post-OE. COBRA triggers tracked in spreadsheets.	Hours of manual data pulls to build a headcount or cost-of-replacement report
With AI	Agent monitors tax changes, notifies you, identifies affected employees, tells you exactly what to update	Agent compiles peer feedback, comp data, and prior notes into a first draft. Manager makes the rating and every judgment call from there.	Continuous eligibility monitoring. COBRA triggers auto-detected. OE drop-off flagged in real time with suggested outreach.	"Show me hires, terminations, and cost of replacement by department for the last 12 months." Dashboard in seconds.

Where AI Creates Compliance Exposure

Where AI Creates Compliance Exposure

AI makes decisions fast, at volume, without calling you first. Scale changes the risk profile.

Benefits Documentation & Notices

- **SPD language:** AI mischaracterizes a plan feature → employee relies on it → denied a benefit → fiduciary liability
- **COBRA notices:** Most litigated benefits compliance issue. AI generates and delivers → employer owns the defect.
- **MHPAEA analysis:** AI-assisted is fine. AI output treated as final without legal review is not.

ACA Reporting & Affordability

- Thresholds change annually. If AI calculates on prior-year figures — and no human verified — every determination may be wrong.
- **1094/1095 filing:** AI-assisted is increasingly common. Has anyone verified the logic reflects current IRS thresholds?

Bias, Fairness & Discrimination Risk

- **Recruiting AI:** Trained on historical data → perpetuates historical bias. Bias audits are a legal risk management practice, not optional.
- **Performance AI:** Reflects whatever patterns — including discriminatory ones — existed in historical review data.

Shadow AI: The Hidden Exposure in Your Own Organization

- **SHRM 2026:** 30% of AI-using workers have violated their org's AI policy. #1 violation: uploading confidential data to an unapproved tool.
- Only 61% ever faced a consequence. The gap between policy and enforcement is where compliance risk lives.

HUB

AI Compliance Scenarios: What Exposure Actually Looks Like

Real exposure. Each of these is happening in organizations running AI-assisted HR and benefits workflows today.

	OE Chatbot Gives Wrong Benefit Information	ACA Affordability Calculation Uses Stale Threshold	Payroll AI Suggests a Correction — No One Reviews It	Recruiting AI Produces a Biased Shortlist
Exposure	Benefit estoppel — the promise AI made may become the plan's obligation	ESRP penalties — $3,340+ per affected full-time employee per year	Wage and hour liability, potential class claims for systematic underpayment	EEOC enforcement, disparate impact liability, potential consent decree
Potential Fix	AI draws only from current, legally reviewed plan documents. Version control + accuracy audits required.	Annual verification that ACA AI logic reflects current IRS thresholds. Document the review.	AI corrections require human approval regardless of amount. Auto-action thresholds must be policy-level decisions.	Independent bias audits. Vendor transparency on training data. Human review before any adverse action.

ERISA Fiduciary Duty Doesn't Transfer to Software

ERISA Fiduciary Duty Doesn't Transfer to Software

ERISA fiduciary duties, act prudently, act in participants' exclusive interest, follow the plan. None transfer to software.

What This Means When AI Is in the Workflow

- **The plan administrator remains the fiduciary.** Delegating a task to AI does not delegate the legal responsibility for the outcome.
- **Vendor selection is a fiduciary act.** Prudent expert standard applies. Evaluate capability, accuracy, data governance, and audit trail, not just features.
- **Monitoring is ongoing.** Accuracy checks, threshold verification, bias audits, these are part of the duty of prudence.
- **Documentation is your defense.** What did you know? What did you do? When? The governance record answers those questions.

Specific Exposure by Area

- **SPD mischaracterization** → Benefit estoppel claim
- **COBRA notice defect** → Up to $110/day per qualified beneficiary
- **ACA reporting error** → ESRP penalties. The AI's error is your penalty.
- **Biased hiring decision** → EEOC disparate impact claim. The AI's logic is not a defense.

The Accountability Chain

- AI vendor produces output → Plan administrator owns the result → Participant has the claim
- The vendor contract may limit vendor liability. The ERISA statute does not limit yours.

The Policy Gap (SHRM 2026)

- **Only 47% of AI-adopting organizations have a governing policy.** No policy means no enforcement. No enforcement means shadow AI. Shadow AI means data uploaded to unapproved tools.

What to Ask Your AI Vendors: A Due Diligence Framework

Ask these at every implementation, renewal, and AI feature release. Treat as a standing checklist.

Question 1 — "What is your AI doing with my employee data?"

- Does it flag, or does it act? Know the difference.
- Is my data used to train your model? If so, is it anonymized? Who owns the outputs?
- **Exposure:** OE chatbots, SPD language, benefits FAQ, enrollment confirmations

Question 2 — “What kinds of decisions is it making, how does it make them, and can I see the logic?"

- Rule-based and auditable, or a model that learns and adapts?
- What happens when AI is uncertain? Does it escalate or produce an output anyway?
- **Exposure:** ACA affordability, eligibility determinations, performance scoring, carrier reconciliation

Question 3 — "When it takes action, what triggers a human review?"

- What actions can AI take without approval? Who defined that threshold?
- Can I produce the audit trail in response to a DOL inquiry or litigation hold?
- **Exposure:** COBRA notices, payroll corrections, carrier feeds, onboarding routing

Non-Negotiables in Every AI Vendor Contract

Signed DPA,
including collection and use limitations, subprocessors, and breach notification

Audit rights,
you must produce an AI output trail; vendor must support it

Model change notification,
advance notice before updates go live in your instance

HUB

What an AI Governance Framework Looks Like

What an AI Governance Framework Looks Like

01 Policy

- Approved use cases and prohibited inputs (PHI, plan data), documented in writing
- Training required before tool use, with clear consequences for violations
- Every AI output touching compliance risk maps to a named human reviewer

02 Vendor Due Diligence

- Written confirmation that inputs don't train the vendor's model
- Signed data processing agreement; legal reviews AI terms at every renewal
- Vendor supports audits of AI-generated output, non-negotiable

03 Human Review Protocols

- No employment decision moves forward without human review
- Benefits-specific: COBRA before transmission, SPDs before publication, ACA thresholds verified annually
- Adoption isn't governance, review must be built into the workflow, logged, and auditable

04 Continuous Optimization

- **Annual inventory:** which tools are active, are outputs still accurate, has the model changed
- Bias and accuracy audits on a schedule, not triggered by a problem
- One owner tracks vendor release notes and vets impact before new features go live

HUB

Building Your AI Governance Infrastructure

From framework to operating practice. Specific assignments, triggers, deliverables.

Step 1

Inventory Your AI Exposure

Written inquiry to every platform vendor: what AI features are active, what data do they use, what can they do without human approval

Deliverable: AI inventory, feature, agency level, data source, current review protocol, owner

Step 2

Tier Your High-Risk Outputs

- **Tier 1, human review before any action:** COBRA, SPD language, ACA determinations, MHPAEA, any output that triggers or denies a benefit
- **Tier 2, human review before delivery:** Enrollment communications, benefits FAQ, onboarding documents, performance review drafts
- **Tier 3, audit log plus periodic review:** Internal analytics, sentiment analysis, retention risk flags, compensation benchmarking

Step 3

Assign Ownership, Not to IT

AI governance in benefits is a fiduciary function. Who reviews the inventory, who approves new use cases, who reviews release notes: define the role

Step 4

Build Into Every Contract Renewal

Require: AI feature disclosure, updated DPA, training data confirmation, audit rights, model change notification

Step 5

Pilot Before You Scale

Start with one low-agency use case. Measure quality. Build the review protocol. Document it. Then expand.

HUB

Using AI to Strengthen Your Team's Capacity

Disruptor or Accelerator? How AI Is Reshaping HR Capacity

Every major technology is both a disruptor and an accelerator. Which one it becomes for your team is a choice.

The Disruptor Side

- Roles focused on routine, transactional HR work are shrinking — now, not eventually
- Employees who don't learn to use AI will have a structural disadvantage against those who do
- **SHRM 2026:** 54% of AI-using workers say pace of work has increased since AI was introduced. 65% say the value of their work has gone up. AI isn't cutting workload — it's reshaping it upward.

The Accelerator Side

- AI handles the admin. People handle the judgment. That's always been where the value is.
- **"How do I use AI to replace the parts of my job I shouldn't still be doing manually?"** That's the right question.
- **SHRM 2026:** AI users report 86% high job engagement — vs. significantly lower for non-users. Engagement follows adoption.

The Critical Thinking Imperative

- **71%** of AI-using workers say critical thinking is more important now because of AI — not less
- Workers who understand AI well use it more cautiously and keep evaluating its output. Those with lower aptitude treat AI as automatically authoritative.
- **Gartner:** "Evaluate how AI reshapes workflows and encourages growth of new roles — not just what process can be automated."

Using AI to Strengthen Your Team's Capacity

Governance without capability is just process. SHRM data shows training approach matters as much as training itself.

For Individual HR & Benefits Contributors

- **Learn to prompt effectively —** specific context produces better outputs. This is a teachable, measurable skill.
- **Know what to review —** threshold dates, plan-specific details, regulatory citations. That's your checklist, not the whole document.
- **SHRM:** Workers below median AI aptitude are 65% more likely to break their org's AI policy — because they can't recognize the compliance risk. Train for aptitude, not just awareness.

For Managers

- **Redefine what "good" looks like —** manager who produces a cost analysis in 15 minutes using AI sets the new standard
- **Build AI review into workflows —** required, logged, with a reviewer who knows what to look for
- **SHRM:** Reward-based training (incentives, competitions, recognition) outperforms punitive approaches by a wide margin. Use carrots, not sticks.

For HR Leaders

- **Assign AI governance ownership —** not IT, not the vendor. The person accountable for the plan.
- **SHRM governance warning:** 55% of Directors & Above would follow AI's recommendation over their own ethical judgment. Build in deliberate practices for independent judgment at the top.
- **Treat AI governance as a fiduciary function —** same rigor as plan admin, vendor selection, and claims oversight

Key Takeaways & Next Steps

Key Takeaways

Six things to take back to your organization today.

Your platforms are already running AI. Ask every vendor — in writing — what's active in your subscription.

The ERISA fiduciary obligation does not transfer to the software. The plan administrator is responsible for AI-generated output. Not the tool. Not the vendor.

Governance gaps are predictable — and provable. SHRM 2026: 30% of AI-using workers have already violated their org's policy. Only 47% of AI-adopting orgs have a governing policy in place.

The SHRM framework gives you a starting point. Policy | Vendor due diligence | Human review protocols | Continuous optimization. Assign an owner today.

AI is both a disruptor and an accelerator. 54% of AI workers say pace of work has increased. 86% report high engagement. Your team's work has already changed.

Train for aptitude, not just awareness. Workers who understand AI follow the rules and apply critical thinking. Workers who don't treat AI as automatically authoritative. That gap is a governance risk.

Next Steps: Your 90-Day AI Governance Starting Point

You don't have to solve everything at once. Do these in order.

Days 1–30

Know What You Have

- Send written AI inquiry to every platform vendor
- Build your AI inventory: feature | agency level | data source | review protocol | owner
- Flag your highest-risk outputs: COBRA | SPD | ACA | MHPAEA | Enrollment confirmations

Days 31–60

Build the Governance Layer

- Draft your AI use policy — approved tools, prohibited inputs, training requirements, consequences
- Assign governance ownership. Name the person. This is a fiduciary function.
- Tier your outputs: Tier 1 (review before action) | Tier 2 (review before delivery) | Tier 3 (audit log)

Days 61–90

Operationalize and Renew

- Review your next vendor contract through an AI lens — DPA, audit rights, model change notification
- Build vendor release note review into your calendar — quarterly minimum
- Schedule your first bias/accuracy audit before next hiring cycle and next plan year

HUB Workforce Solutions can help with

AI vendor due diligence | Governance policy development | HR technology strategy | Ongoing compliance oversight

Q&A

Upcoming Events

hubinternational.com/events/

WEBINAR – Employee Benefits

Women's Health & Wellbeing at Work:

What Employers Can Do About Support, Policy and Culture

Wednesday, October 21 | 11 AM CT

Thank you

For more information visit **www.hubinternational.com**

HUB

APPENDIX

Glossary of Terms

- **ACA** — Affordable Care Act; federal law requiring applicable large employers to offer affordable, minimum-value health coverage or face penalties
- **Agentic AI** — AI that takes a goal, executes the necessary steps, and follows up without being re-prompted
- **Benefit Estoppel** — A legal claim that a plan may be bound to honor a benefit an employee reasonably relied on, even if it conflicts with the written plan terms
- **COBRA** — Consolidated Omnibus Budget Reconciliation Act; federal law allowing employees to continue group health coverage after a qualifying event
- **Conversational AI** — Chat-style AI that answers questions in plain language, typically drawing from approved sources
- **DOL** — Department of Labor; federal agency responsible for enforcing labor and benefits laws, including ERISA
- **DPA** — Data Processing Agreement; a contract governing how a vendor may collect, use, and protect an employer's data
- **EEOC** — Equal Employment Opportunity Commission; federal agency enforcing laws against workplace discrimination
- **ERISA** — Employee Retirement Income Security Act; federal law setting minimum standards for retirement and health plans, including fiduciary duties
- **ESRP** — Employer Shared Responsibility Payment; the ACA penalty assessed on employers that fail to offer required health coverage
- **Fiduciary Duty** — The ERISA obligation to act prudently and solely in participants' interest; the duty belongs to the plan administrator, not the software they use
- **HCM** — Human Capital Management; software used to manage HR functions such as payroll, benefits, and workforce data
- **IRS** — Internal Revenue Service; federal tax agency responsible for ACA reporting requirements and annual thresholds
- **MHPAEA** — Mental Health Parity and Addiction Equity Act; federal law requiring parity between mental health/substance use and medical/surgical benefits
- **Multimodal AI** — AI that works across text, voice, images, documents, and video within the same tool
- **OE** — Open Enrollment; the annual period when employees select or change their benefit elections
- **Predictive Analytics** — AI that draws on historical data to flag trends, such as turnover risk or utilization, before they become problems
- **Shadow AI** — Unapproved or unsanctioned use of AI tools by employees, outside of IT or compliance oversight
- **SHRM** — Society for Human Resource Management; professional HR association and source of the 2026 AI workplace research cited in this webinar
- **SPD** — Summary Plan Description; the required participant-facing document summarizing a benefit plan's terms