

The cover image is split into two main vertical sections. The left section is a solid orange color with a faint, pixelated world map in the background. The right section is dark blue with vertical streaks of light in shades of blue, cyan, and red, along with small circular bokeh effects. The title 'AI Briefing for Industry Leaders' is written in white, bold, sans-serif font on the orange background.

AI Briefing for Industry Leaders

BrightLine
Solutions

SEPTEMBER 2025

The AI Paradox Point

Dear Reader

First off, I want to thank you for your feedback after last month's issue, and for your requests for the specific topics you want us to explore.

Consider it done: these topics have found their way into this month's briefing.

Also, big thank you for a whopping **38% open rate** with just one unsubscribe!

Last minute special section: On September 26, 2025, the Government of Canada launched the [AI Strategy Task Force](#) and a 30-day national engagement sprint to shape “the next AI strategy” with themes spanning infrastructure, industry adoption, skills, safety, and investment. Expect a refreshed federal AI playbook to follow the sprint. BrightLine Solutions will be covering this task force and national AI strategy, and what it means for you, starting this month.

This last week I had the honour of moderating a cybersecurity panel for 100+ association executive directors, board chairs and other executive leaders, as part of the inaugural issue of [talkEXEC](#) in Fredericton, NB. Unsurprisingly, the first question from the audience was about the use of AI in cyberattacks.

Well, two days later, [Noma Labs announced](#) the “ForceLeak” incident involving an expired domain and indirect prompt injection attacks targeting AI processing agents. Using this technique, hackers can leverage your CRM intake forms’ description field to secretly embed malicious commands inside your customer data, which are eventually picked by your CRM’s internal AI agent (like Salesforce’s AgentForce), which mistakes the code for a legitimate task and causes it to, for example, autonomously exfiltrate sensitive CRM records without triggering security alerts.

But panic not: Salesforce has since patched this vulnerability. Just take this as a lesson to check third party written prompts before feeding them to your AI agents!

What this month's research revealed, are the peculiar **paradoxes of AI adoption**, and that is what we decided to focus on.

Consider this: a CFO just approved another million-dollar investment in AI, marking her department's third such investment this quarter. The team is now part of 72% of finance departments using AI in 2025, up from 34% in 2024. At the same time, the same CFO is supporting a petition for increased government regulation of AI technology. This story is fictional, but it is based on [real statistics](#).

This is the defining contradiction of our moment: we're simultaneously racing toward AI adoption and worshipping it, while desperately trying to pump the brakes on its grow. And both impulses are correct.

Think of what happened earlier this month in a San Francisco court. [Anthropic agreed to pay \\$1.5 billion](#), the largest copyright settlement in U.S. history, to authors who alleged that nearly half a million of their books had been illegally pirated to train Anthropic's models.

Meanwhile, in warehouses across Britain, [DHL is deploying an army of 1,000 new robots alongside their existing 2,000 mechanical workers](#). Each of these "Boston Dynamics Stretch" robots can unload 700 boxes per hour, a rate that would break any human's back by lunch time. DHL frames this as "collaborative robotics" machines working alongside people. But if you ask the 83% of workers who fear AI will eliminate entry-level jobs, they'll tell you a different story about collaboration.

As we stand at this **paradox point** in which we both embrace and fear AI, let's continue to engage thoughtfully, advocate for responsible innovation, and shape the future together.

And finally, I invite you to [email me](#) your perspectives on these topics, and your preference for future issues, so we can collectively navigate the complexities of AI adoption in our organizations and communities.

Welcome to September's AI Briefing!

A handwritten signature in black ink that reads "Nestor Gomez". The signature is written in a cursive, flowing style.

Contents

Finance: The Believers and the Burned5

HR: The Skills Gap Nobody Discusses5

Legal: When Machines Learn to Lie6

Operations: The Robot Warehouse Reality7

Security: The Attackers Have AI Too7

Special: Canada’s AI Task Force.....7

 What's Actually Happening8

 What We Have Noticed Thus Far.....8

 What Matters the Most for Your Planning Horizon8

 What to Monitor in Q4 20258

What All This Means for Your Organization9

Finance: The Believers and the Burned

The [Protiviti's Global Finance Trends Survey](#) results landed on my desk with a thud of inevitability. Finance departments have become true believers and 72% now use AI tools daily, treating spreadsheets like ancient scrolls.

Christopher Wright, who leads Protiviti's CFO practice, puts it bluntly: "CFOs are no longer simply stewards of capital who report results".

But here's where it gets interesting. While CFOs are all-in on AI, they're also the first to spot the cracks. [MIT's Project NANDA discovered that 95% of generative AI pilots fail to deliver rapid revenue acceleration](#). That's a human problem.

The pattern is telling: Purchased AI tools succeed two-thirds of the time, while internally built tools succeed only one-third of the time.

Why? Because most organizations try to boil the ocean instead of heating a cup of tea. The winners **focus on one specific pain point, empower the people doing the work, and resist the urge to revolutionize everything at once.**

Then [Visa dropped a bombshell](#). In April, they announced AI agents that can spend your money autonomously. Not only suggest purchases, but complete them.

Giving away our control. That's the biggest differentiation factor of this technological revolution compared to digitization, electricity, steam, etc.

The system is live now, with partnerships spanning OpenAI, Anthropic, Microsoft, and others. Your chatbot can now book your flight, order your groceries, and max out your credit card, all while you sleep.

The financial industry calls this progress. Security experts call it a disaster waiting to happen. A paradox.

[Commonwealth Bank of Australia](#), for example, just handed ChatGPT Enterprise licenses to all 52,000 employees. [Wells Fargo partnered with Google](#). [Santander went all-in with their own AI initiatives](#). The message is: embrace AI or become irrelevant. Yet each partnership opens new attack vectors that didn't exist six months ago.

HR: The Skills Gap Nobody Discusses

A new survey from General Assembly shows that only [22% of business leaders believe entry-level staff are "fully prepared" for their jobs](#). Interestingly, this is a 10% perception increase from 2024's survey.

Meanwhile, the same source cites that 83% of workers fear AI will replace those same entry-level positions.

Are we simultaneously claiming young workers aren't ready while building systems to eliminate their jobs entirely. What does this mean for young graduate professionals?

The disconnect runs deeper. A [Paychex survey](#) found 72% of small business owners are optimistic about AI, with 65% already using it daily. Yet only 53% of HR leaders plan to invest in AI for recruitment this year. The very department responsible for preparing the workforce for an AI future is dragging its feet on adoption.

The [McKinsey survey](#) is large and, while US-weighted, its contradictions matter: employees rank formal training as the single most important lever for adoption, yet nearly half say they receive only moderate or less support and more than one in five report minimal to no organizational support. People are moving faster than their bosses think, they want help, and many employers are still leaving them to figure it out.

Executives, meanwhile, promise cash for the AI future, with 92 percent saying they plan to boost AI investment over the next three years, yet only 1 percent describe their organizations as "AI mature".

That mismatch explains why we see a flurry of pilots and platform buys that rarely translate into scaled value: leaders are spending on technology, while the workforce asks for training and operational integration.

The practical conclusion is that, if boards want ROI, they must stop treating training as an HR nice-to-have and start funding role-specific, hands-on capability programs tied to measurable workflow outcomes, otherwise the dollars poured into tech will mostly wallpaper over a leadership problem.

Legal: When Machines Learn to Lie

The [Anthropic settlement](#) tells us everything about how we're handling AI's legal challenges. Which is to say, badly.

Judge William Alsup ruled that training AI on copyrighted material is "transformative" fair use. His reasoning is that AI models are like "any reader aspiring to be a writer", they learn from existing works to create something new.

This anthropomorphization (try that word instead of acetaminophen!) of machines would be quaint if it weren't so dangerous, because machines don't "aspire". They process probability distributions.

[Thomson Reuters found](#) that 80% of legal professionals believe AI will have "transformational" impact within five years. They're already seeing returns: 59% report time savings of up to 240 hours annually. But here's what they're not saying: those time savings come from automating the entry-level work that traditionally trained young lawyers. We're pulling up the ladder behind us.

The Canadian statistics are even more telling. The Leger poll found [85% of Canadians want AI regulation](#). Not guidelines but hardcore regulation. Interestingly, comments from AI Minister, Evan Solomon, suggest the government is focusing on adoption instead. The public demands guardrails while government hits the accelerator. How do we balance that?

Operations: The Robot Warehouse Reality

[DHL's announcement](#) reads like science fiction: £550 million to deploy 1,000 robots by 2030, adding to their existing army of 2,000 mechanical workers. The Stretch robots from Boston Dynamics work tirelessly, unloading containers at superhuman speed.

But visit one of these warehouses, and you'll see humans everywhere. They're not packing boxes anymore; robots do that now. They're not driving forklifts; robots do that too.

Instead, humans troubleshoot when robots fail, make decisions robots can't, and provide the judgment that no algorithm can replicate. Yet.

The company insists this is about "collaboration", and that robots reduce physical strain on workers.

But is hard to strain your back when you don't have a job, right? The reality is that every efficiency gain, every productivity improvement, and every cost reduction requires humans to adapt, and this is what's usually absent on the balance sheet.

Security: The Attackers Have AI Too

Remember when phishing emails were easy to spot? The broken English, the obvious scams?

Reuters just demonstrated that ChatGPT and other leading chatbots can craft phishing emails indistinguishable from legitimate communication. Worse, they can personalize these attacks at scale, targeting thousands of victims with customized deception.

The Salesloft breach in August should have been our wake-up call. Attackers exploited AI-driven sales chatbots to steal Salesforce tokens, exposing entire customer databases. The irony is that the tools we built to enhance security became the very weapons used against us. Not the first time, is it?

In the UK, 35% of firms reported AI-powered fraud attempts in just the first quarter of 2025. These aren't Nigerian princes, but sophisticated operations using voice cloning, deepfakes, and behavioral analysis to exploit human psychology.

But let's close with a positive note here: the UK government's Fraud Risk Assessment Accelerator did manage to recover £480 million using AI detection systems. Yay!

Special: Canada's AI Task Force

Evan Solomon, Ottawa's minister of AI and Digital Innovation [announced an AI Strategy Task Force](#) to inform Canada's next national AI strategy

Included is a 30-day consultation sprint through October. For executives in food and beverage and association management, this deserves careful attention with cautious optimism.

What's Actually Happening

The federal government decided to refresh the [Canada's 2017 Pan-Canadian AI Strategy](#) with input from 30 task force members spanning academia, industry, and civil society, organized in eight working groups: Research and Talent; Adoption (industry and government); Commercialization; Scaling Canadian AI and Attracting Investments; Safe AI; Education and Skills; Infrastructure; Security.

What We Have Noticed Thus Far

- Public consultation runs from October 1 to 31, with recommendations expected in November.
- The task force includes heavyweight names like Patrick Pichette from Inovia Capital, Louis Têtu from Coveo, and academics from leading AI institutes.
- **The task force's composition includes tech industry, venture capital, and academic interests, but misses voices from operational leaders in agri-food, associations, and non-profits.**
- There is limited representation from mid-market operators who are most affected by the real-world challenges of AI adoption.
- The 30-day sprint for stakeholder engagement matches the federal government's urgency to establish Canada as a global AI leader.

What Matters the Most for Your Planning Horizon

1. **Adoption stream:** co-chaired by Dan Debow and Cari Covent, this track could accelerate AI procurement frameworks that influence how organizations engage with government and procure Canadian AI providers.
2. **Infrastructure:** led by Garth Gibson, this workgroup may determine whether SMEs gain affordable access to compute and cloud resources currently dominated by large enterprises.
3. **Education and skills group:** including First Nations Technology Council's Natiea Vinson, could shape workforce development programs relevant to your talent and change management efforts.

What to Monitor in Q4 2025

Watch these three critical developments:

1. **Task force recommendations.** Monitor if the recommendations address the pilot-to-production gap that traps 95% of generative AI projects.
2. **Infrastructure proposals.** Track if these include regional capacity accessible to organizations outside major tech hubs.

-
3. **Adoption.** Whether the guidelines move past generic frameworks and into sector-specific pathways for agri-food, safety systems, member engagement platforms, and operational workflows.

Consider submitting feedback on barriers, data sovereignty concerns, or any other topic you think must be addressed.

As of the date of this publication, Sep 29, 2025, there's no clear intake mechanisms for contributing with the task force. We recommend to watch for further announcements on the **two official paths** and **one invited path** announced.

1. **Watch for the "Consulting Canadians" portal (official).** ISED's announcement says Canadians are invited to share perspectives through the government's Consulting Canadians portal during the 30-day sprint. Watch for a dedicated consultation page to go live there and submit your input.
2. **Respond to the consultation once it's posted under ISED's public consultations.** ISED lists [active consultations](#) on its site, and I expect the AI strategy item should appear there with instructions.
3. **Task Force members' outreach (invited networks).** The government says Task Force members will consult "their networks" for actionable recommendations. If you're connected to any of the named leaders through industry, academia, associations or other organized bodies, share a short brief directly through those channels.

What All This Means for Your Organization

Begin with a candid evaluation: what challenges your organization is facing, and which ones are simply the product pitches of generative AI vendors? That gap between what you need and what you're being sold is where resources are wasted and potential is lost.

Don't get distracted by hype or the headline-grabbing projects. Returns come from focusing on the essential, everyday tasks such as automating invoices, reviewing contracts, or managing routine customer requests. These are the practical areas where generative AI delivers steady value right now.

Make your people your primary investment, not despite AI, but because of it. Organizations making AI work are defined by people who know when to trust it, when to intervene, and when to ask the tough questions.

Accept the paradox. Staying competitive means adopting AI, but it also means facing the existential threats it brings. Both realities exist side by side. Ignoring this duality is a sure path to failure.

Use the national AI strategy consultation as a lever. Stay informed on the AI Strategy Task Force's intake consultation channels and ask what will move your numbers: adoption vouchers for SMEs, compute credits to offset cloud costs, simpler procurement, and sector training. Tie your ask to a concrete KPIs and commit to match funding or publish outcomes.

We're witnessing the most profound technological transformation in decades. There's no roadmap, no safety rails, and no certainty about the destination...yet.

The only thing we can count on is that the old playbook is obsolete.

Your CFO will continue funding AI initiatives while simultaneously calling for stricter regulations and asking for a business case.

HR will lean into automation for recruitment yet mourn the diminishing role of human judgement.

Operations will deploy more robots, even as they assure workers their jobs are safe.

This is the new normal, not hypocrisy. The sooner we stop pretending these contradictions can be resolved, the sooner we'll learn to operate within them.

Each solution breeds new challenges. Every efficiency new complexities. Every step forward uncovers fresh risks.

We've constructed a paradox engine, and now we're all inside it. The challenge isn't to escape, but to survive and thrive while the gears keep turning.