An Enterspeed white paper

10 things you need to know about AI and governance

By CEO Toke Lund

www.enterspeed.com

Contents

Intro

Let me just say it. Most retailers are doing AI governance backwards.

They build compliance frameworks before they understand what they're actually trying to control. Nobody would install a sophisticated alarm system before figuring out which doors need locks. The same logic applies to AI governance – and yet here we are.

Most retail organisations now use AI in at least one business function. Generating product descriptions. Translating content across markets. Personalising campaigns. But a significant number feel overwhelmed by the regulatory and operational reality of governing it. Leaders expect ROI within one to three years. What they don't always plan for is the infrastructure work required to get there safely – without a hallucinated product spec, a GDPR violation, or a brand inconsistency reaching a customer.

So here is what you actually need to know about AI governance. Not the theoretical framework – the real challenges that will make or break your AI initiatives in retail.

Enterspeed CEO, Toke Lund

1

Governance isn't optional. It's infrastructure

Governance doesn't slow down AI adoption. Poor governance does.

AI governance has evolved from an operational consideration to a fundamental requirement. Think of governance the same way you think of your data infrastructure. It's not something you bolt on after the fact when regulators come knocking – or when a product page goes live with the wrong price in three markets.

For retail, the stakes are immediate. Wrong information reaches customers. Brand inconsistencies compound across channels. Compliance failures aren't abstract – they show up on product pages, in campaign copy, and in customer-facing communications.

Companies winning at this treat governance as a technical capability, not just a compliance checkbox. Organisations already certified under ISO 27001 are well placed to adopt ISO 42001, the AI management standard, since many controls overlap. If your data house is already in order, you're further along than you think.

“Think of governance like you think of your data infrastructure. It's not something you bolt on after the fact when regulators come knocking.”

2

Human oversight must evolve: From "in-the-loop" to "*on*-the-loop"

You can't approve every AI output. But you can build systems that tell you which ones need your attention.

The traditional model – where humans review every AI output – is collapsing under its own weight. The volume of AI-generated content in a modern retail operation makes one-to-one review impossible.

The answer is a shift to "human-on-the-loop": AI handles scale and flags uncertain or high-stakes cases for human review. You maintain ethical oversight without creating a bottleneck that kills efficiency. For retail specifically, this means your compliance-sensitive product claims, legal text, and pricing information get human eyes – while routine copy flows through.

3

Control your data or AI will control your reputation

Your AI is only as trustworthy as the data governance underneath it.

AI systems depend on data, but uncontrolled data flow is a governance nightmare.

One of the biggest risks is what you might call "oversharing" – AI tools surfacing sensitive data to the wrong people, or worse, to customers. With GDPR and the EU AI Act now in force, the consequences are significant: fines of up to €35 million or 7% of global turnover for serious non-compliance.

Your data supply chain must respect permissions natively. If your marketing team shouldn't see finance data, your AI shouldn't either. This isn't a technical detail – it's a business boundary.

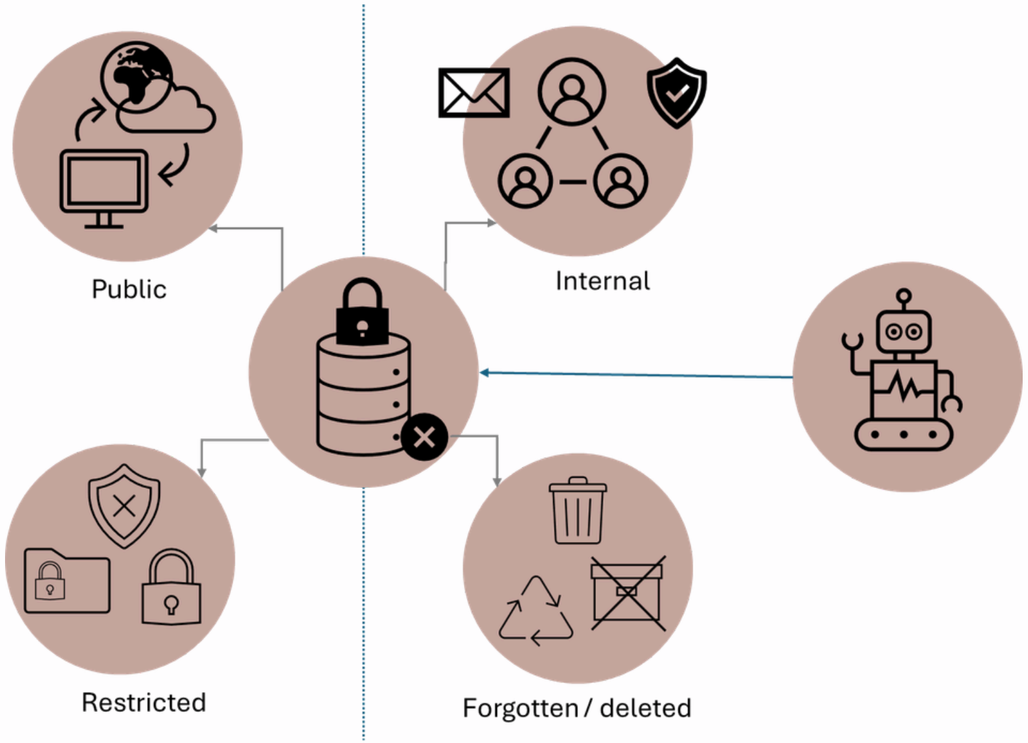

4

Context is king – without it, AI hallucinates

Context isn't a nice-to-have. It's what separates useful AI from liability.

Large language models are probability machines. When they lack knowledge, they fill gaps with plausible-sounding nonsense.

For a retailer, that means product descriptions that invent specifications, pricing copy that contradicts your actual offers, or compliance text that doesn't reflect current regulations.

The fix is contextual grounding – feeding models domain-specific, real-time information rather than relying on what they "remember" from training. Protocols like the Model Context Protocol (MCP) standardise how AI connects to your data.

Without this contextual grounding, you are just asking AI to read your company's story with half the pages missing.

“Without this contextual grounding, you are just asking AI to read your company's story with half the pages missing.”

5

Decoupled architecture scales better and is more secure

Flexibility and control aren't opposites. The right architecture gives you both.

To govern AI effectively, you need an architecture that can flex. A decoupled approach separates the frontend presentation layer from the backend data and logic, making it far easier to enforce security and compliance policies through APIs rather than rewriting entire systems.

In practice, that means your product catalogue can serve your webshop, your app, and your AI layer independently – without one change breaking another.

If you're building monolithic structures today, you're building technical debt for tomorrow. Decoupled architecture isn't just a technical preference – it's a governance enabler.

See next page for a visualisation of this.

Monolithic (before)

One big system

- Webshop frontend
- Mobile app
- AI content layer
- Product catalogue
- Compliance rules

One change risks breaking everything

Decoupled (after)

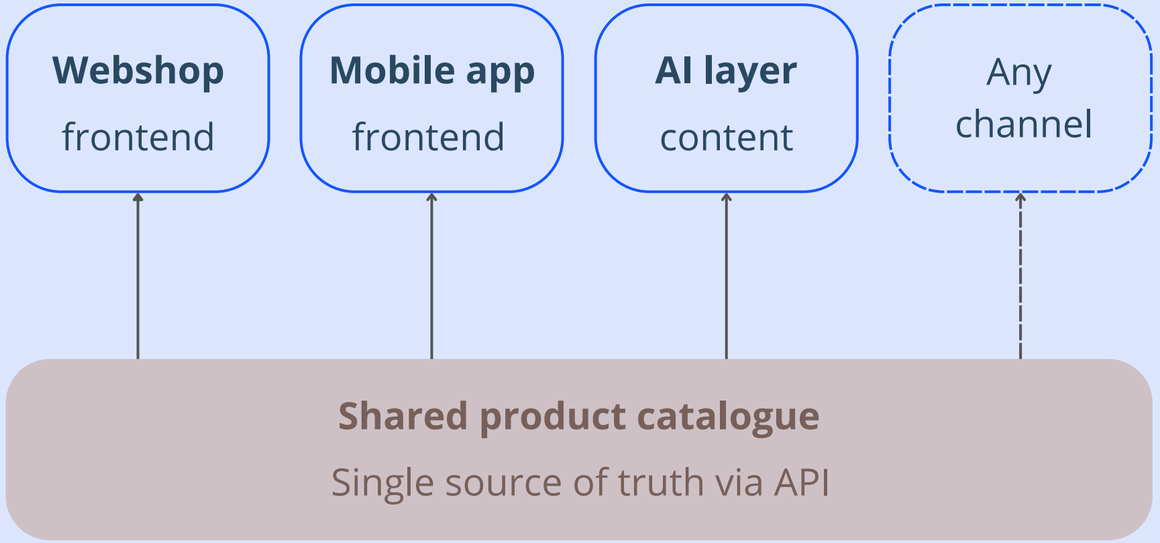

Each layer works independently

6

Your AI is only as good as your data supply chain

Dark data isn't just wasted potential. In an AI-first environment, it's a risk.

A significant portion of enterprise data is never actually used – stored, forgotten, and completely invisible to the people trying to govern it. The moment AI systems can query that data, it becomes an enormous governance and security surface area.

Your AI doesn't just need data. It needs the right data, at the right time, with the right context. That's a supply chain problem. You need lineage – where did this data come from? – and transformation logic – what happened to it along the way? Without this, even the most advanced models will generate inaccurate, untrustworthy outputs.

“Your AI doesn't just need data.
It needs the right data,
at the right time,
with the right context.”

7

Transparency builds trust

Trust isn't built by hiding how AI works. It's built by being able to show your work.

Fewer than half of people globally say they trust AI systems – and in a retail context, that distrust translates directly into customer hesitation and internal resistance to adoption.

Transparency is the answer. Not exposing every neural connection, but making the data flow visible and auditable. Documenting where training data came from, logging when humans intervened, creating a trail that can be followed when something goes wrong. This isn't just about keeping auditors happy. It's about building systems you can actually debug – and defend.

8

Accountability cannot be automated

When things go wrong, "the AI did it" isn't a defence.

Even as AI becomes more autonomous, humans remain accountable. That's not a limitation – it's a safeguard. The temptation might grow to treat it as an independent actor, but it isn't. Courts and regulators will not accept "the algorithm decided" as a defence.

Ethical responsibility for AI systems remains a human obligation. Machines can track activity but humans own the outcomes.

This means clear roles, clear escalation paths, and a culture where accountability doesn't get quietly offloaded to a model.

“Even as AI becomes autonomous, humans remain accountable.”

9

Privacy by design, not by accident

Privacy isn't a feature you add. It's a foundation you build on.

Data privacy is a top concern for the majority of organisations running AI at scale – and yet privacy is still frequently treated as an afterthought, bolted on once the system is already built.

Privacy by design means structuring your data architecture so that protections are inherent from the start. It means minimising data exposure, ensuring that AI training data is handled with the same rigour as production data, and not inadvertently multiplying your attack surface by copying sensitive datasets into new environments.

10

The goal isn't perfect AI – it's responsible AI

Trustworthy AI isn't about having perfect models. It's about having the right systems in place when they're imperfect. And they will always be imperfect.

Perfect AI doesn't exist. What exists is AI that is good enough for your use case, surrounded by guardrails that catch the errors before they cause damage. The companies getting this right aren't waiting for perfection – they're adopting recognised frameworks, building feedback loops, and improving continuously.

The retailers moving fastest with AI aren't the ones who've solved governance once and moved on. They're the ones who've built governance into how they operate – as infrastructure, not overhead.

Get your ducks in a row

So, what's the (uncomfortable) truth?

What I've learned building data supply chains is that most companies struggling with AI governance actually have a data problem, not a governance problem.

They try to govern data they don't fully control. Contextualise data they don't fully understand. Secure data they can't fully track. Fix your data supply chain first – and AI governance becomes considerably simpler, because you're no longer trying to impose order on chaos.

Start with three questions: Where is your AI operating without a reliable, real-time data source? Which use cases carry the highest business risk if something goes wrong? And do you have visibility into what your AI is actually doing?

If you're not sure where to start, we'd be happy to walk you through how other retailers are approaching this.

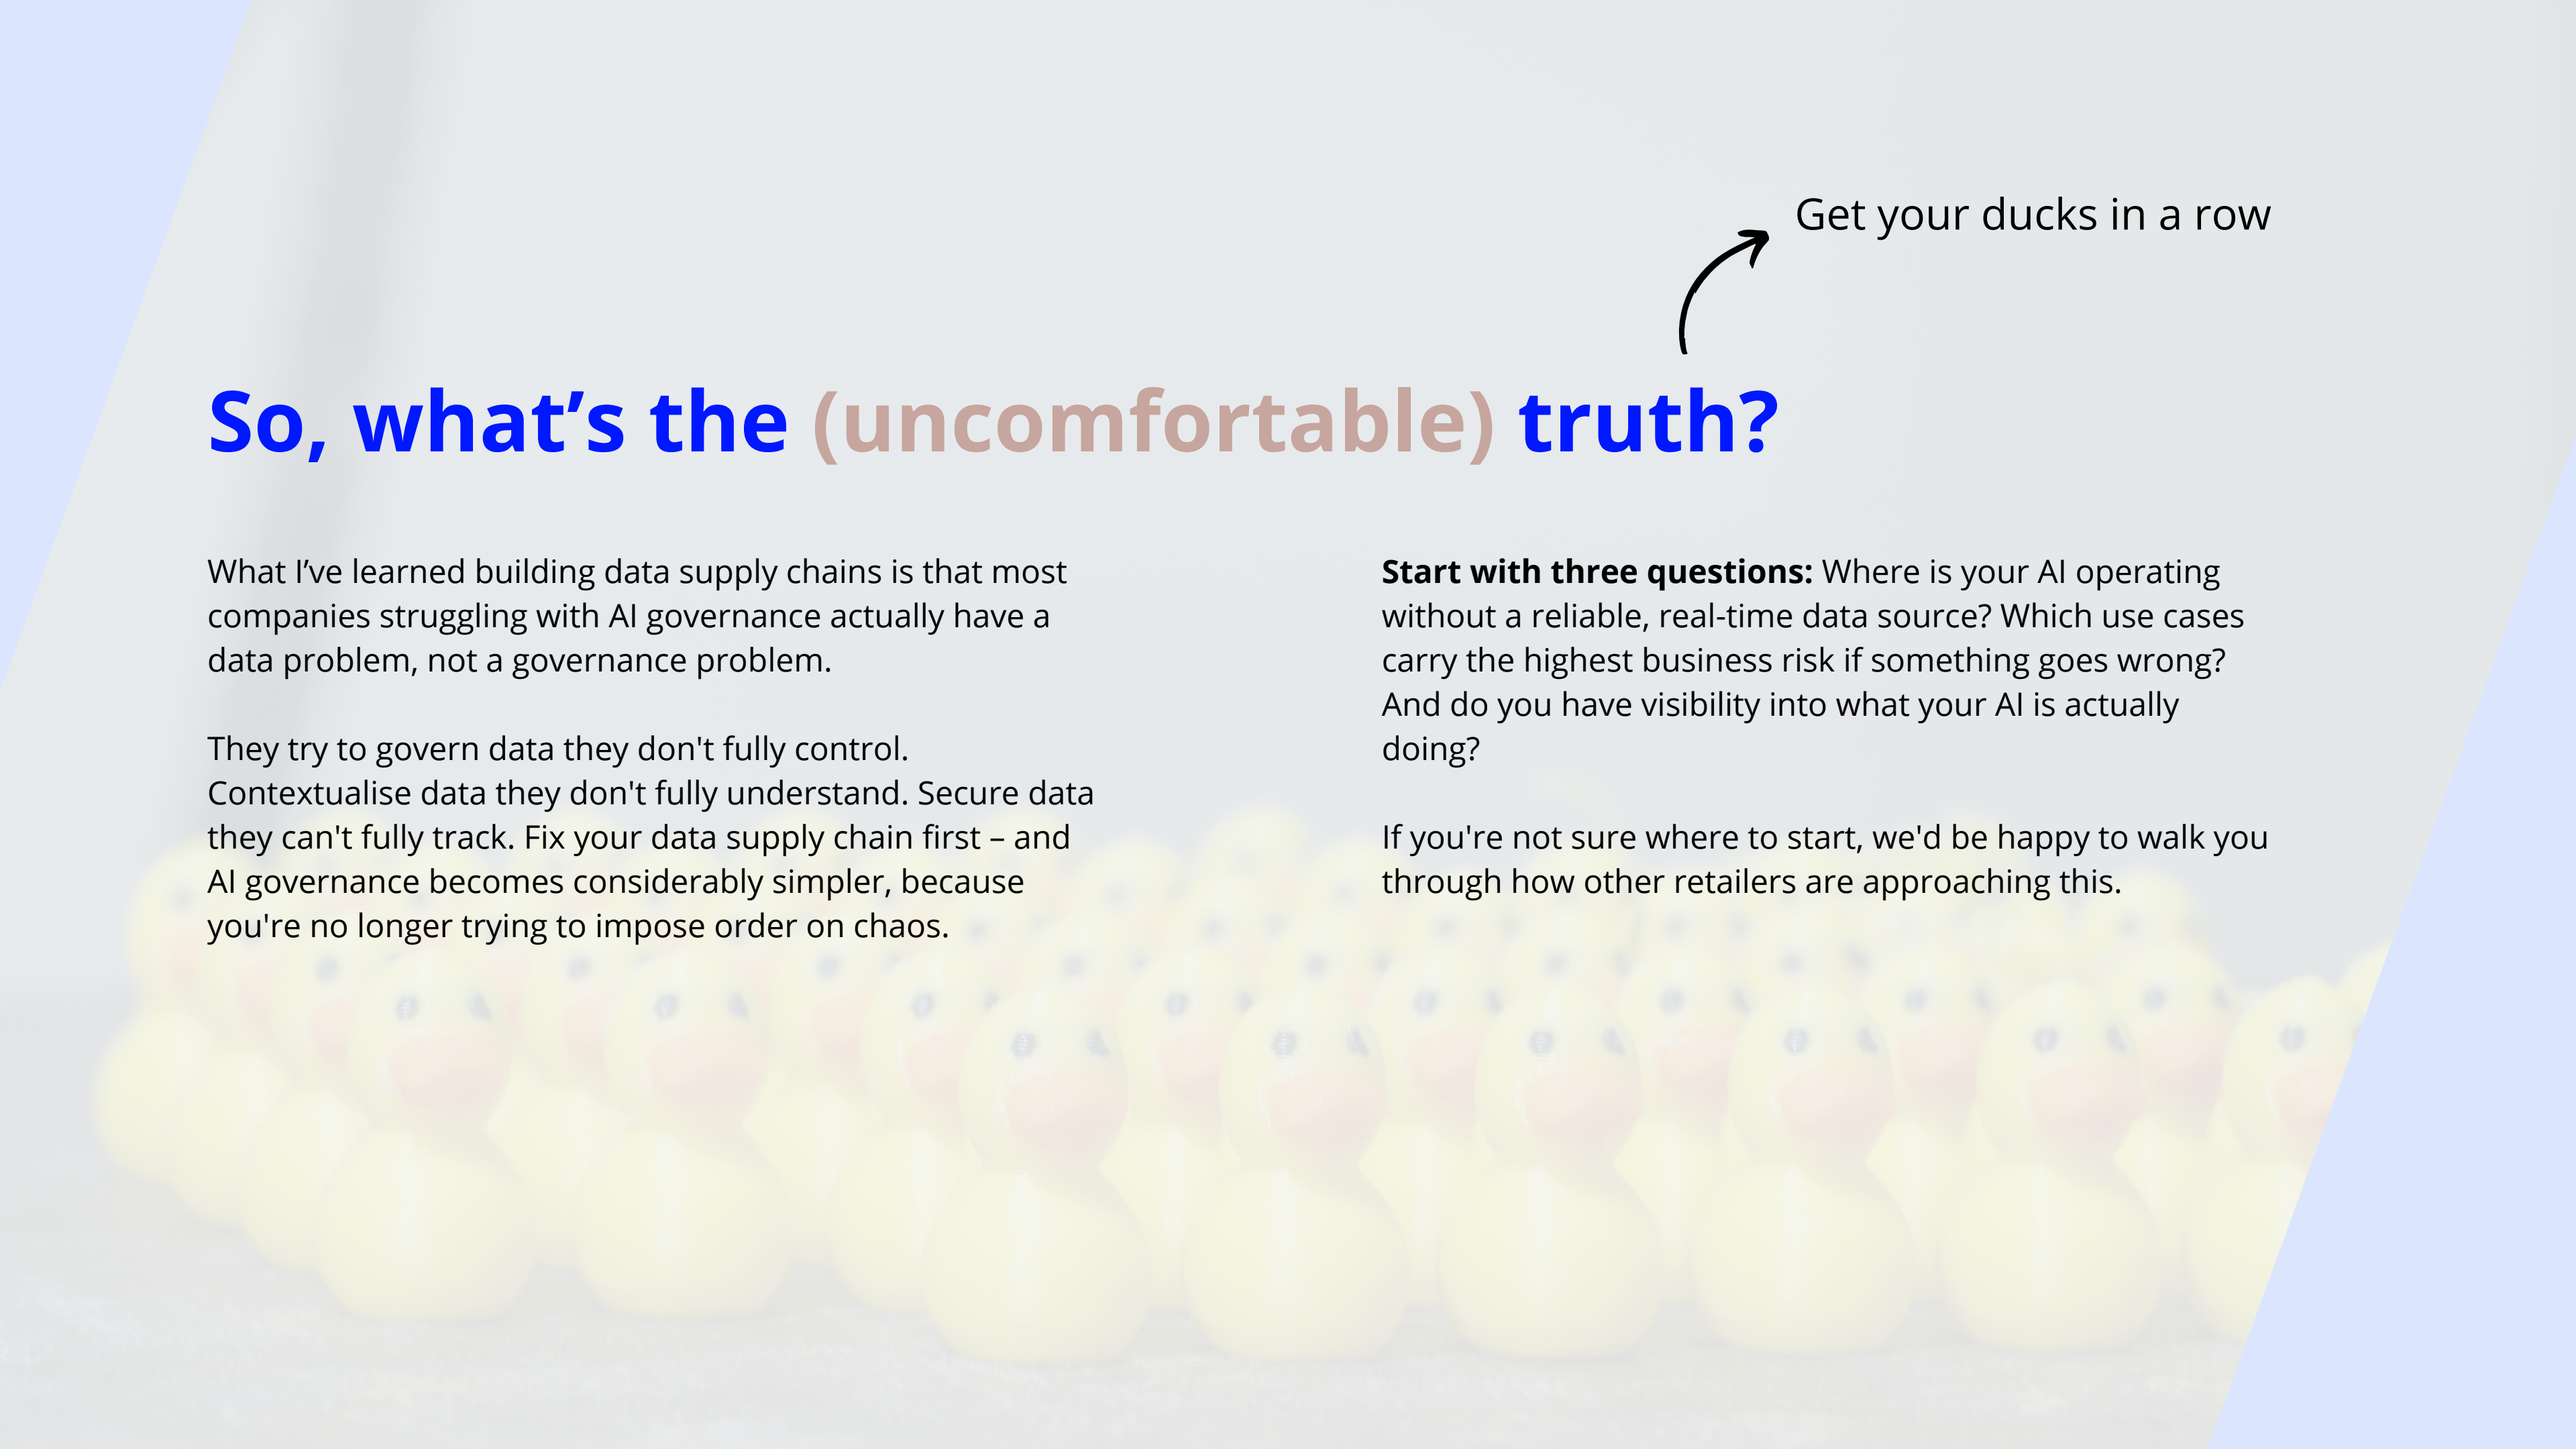

About Enterspeed

Enterspeed accelerates innovation by minimising complexity.

We bridge new and legacy systems, and give enterprises a single, trusted data supply chain for AI, digital orchestration, and agentic commerce.

Trusted by enterprises across Denmark, Norway, and the UK.

Our mission? Accelerate growth. Reduce complexity.

We enrich and structure your data delivery
and break down silos to ensure a single source of truth for GenAI and digital operations.

We decouple and abstract systems
eliminating dependencies on rigid architectures and providing full flexibility.

We optimise for GenAI
by delivering lean, structured, and real-time data to any model.

We reduce costs and complexity
ensuring lower operational costs by minimising maintenance, infrastructure overhead, and development effort.

And above all... **we empower digital teams**
with a simple, API-driven approach that enhances speed, scalability, and ease of integration.

Enterspeed

Accelerate growth. Reduce complexity.

Get in touch with us anytime
info@enterspeed.com
+ 45 21 55 18 25

Enterspeed A/S, P.O. Pedersens Vej 2, DK-8200 Aarhus N, Denmark