Skyjed

WHITEPAPER
AWS SUMMIT SYDNEY · MAY 2026

You can now build at 15–20× speed.

Does your product lifecycle governance keep up?

A CTO's honest reckoning with what AI-driven development changes — and the lifecycle visibility, oversight and governance guardrails required to iterate for growth at AI speed, compliantly and at scale.

Stephen Brown
Co-founder & CTO, Skyjed

RELEASED
AWS SUMMIT SYDNEY
MAY 2026

CONTENTS

In this paper

The speed problem is solved. The governance problem is just beginning. This paper names the gap, explains why existing frameworks weren't built for it, and describes what governance must look like to operate at AI velocity.

A NOTE FROM THE AUTHOR

I didn't set out to write a paper about governance.

I came to AWS Summit as a CTO who had just done something that genuinely surprised me. My team had compressed a feature build we'd estimated at four to six weeks into 48 hours. Not a prototype. Not a demo. Production-ready code, shipped to enterprise clients. And I wanted to understand why it worked — and what it meant.

The methodology was AI-DLC. The tool was Kiro. The result was real.

But the more I thought about what we'd built, the more a different question started forming. We'd solved the speed problem. We'd proven the methodology. And then I walked back into our own platform — a product lifecycle governance platform, of all things — and I had to ask myself: does our product lifecycle governance keep up with what we just demonstrated we can build?

This paper is my attempt to answer that question honestly. Not as a vendor. As a CTO who has lived on both sides of the equation — the speed, and the discipline required to govern what that speed produces.

If you've just come from a session celebrating what AI-driven development makes possible, this paper is the next conversation. The one that happens when the enthusiasm settles — and you realise that to truly iterate for growth, you need the lifecycle visibility to see what you've built, the oversight to govern it, and the governance guardrails to do it with confidence.

— Stephen Brown

Co-founder & CTO, Skyjed

EXECUTIVE SUMMARY

Speed is solved. Governance is the next strategic problem.

Something fundamental shifted in enterprise software development in the last twelve months. Teams adopting AI-driven development are reporting 15–20× acceleration in feature delivery. The evidence is no longer anecdotal. It is in production, at scale, in regulated enterprises.

15–20× faster feature delivery

Cross-functional teams using the AI-Driven Development Lifecycle (AI-DLC) are compressing what was four-to-six-week feature builds into 48 hours of production-ready code.

This is genuinely good news. Faster delivery means faster value. Faster value means competitive advantage. The organisations that master AI-driven development will pull ahead of those that don't — and the gap will compound quickly.

But speed creates a second-order problem that no one on the conference stage is talking about yet.

Every product and asset an organisation ships carries lifecycle obligations from inception to sunsetting. Regulatory alignment. Strategic coherence. Performance visibility. Change and decommission readiness. These obligations existed before AI-driven development, and they persist after it. **What has changed is the rate at which they accumulate.**

When development was measured in weeks, governance had time to keep pace. Stage gates, risk reviews, and portfolio assessments were slow — but development was slower. The system, however imperfect, stayed roughly in balance.

At AI speed, that balance breaks. Assets ship before governance teams have been briefed. Portfolios grow faster than they can be assessed. Regulatory exposure accumulates in the gap between what was built and what was reviewed. Signals get lost. Lifecycle drift sets in quietly. And in regulated industries — financial services, insurance, energy — that gap is not an operational inconvenience. **It is a material risk.**

DEVELOPMENT HALF

- AI-DLC compresses delivery **15–20×**
- Cross-functional teams ship in days, not weeks
- Quality validated by collective context
- Adoption is spreading across regulated enterprises

GOVERNANCE HALF

- Stage gates calibrated to weekly cycles
- Quarterly portfolio reviews, annual risk assessments
- Manual evidence collection via spreadsheets
- Governance designed for a slower world — left in place

The thesis

This paper names that risk. It explains why existing governance frameworks were not designed for AI-speed delivery. And it describes what product lifecycle governance must look like to move organisations from reactive to proactive — with the right visibility, the right oversight, and the right human judgment at the moments that matter.

> *The organisations that win in the AI era will not be the fastest builders. They will be the ones who build fast and govern well — turning AI-speed development into durable competitive advantage, compliantly and at scale.*
>
> — STEPHEN BROWN, CTO, SKYJED

PART 01 THE SPEED IS REAL

And that's not the problem.

What I learned in 48 hours of AI-DLC, why structure — not the AI — was the unlock, and why coming back to our own platform changed the question I needed to answer.

PART 01 · THE SPEED IS REAL

What we built in 48 hours

I'll start with what actually happened, because the most useful thing I can do is tell an honest story rather than a polished one. Earlier this year I participated in a two-day AWS-guided experiential learning session on the AI-Driven Development Lifecycle. I brought a real problem — a feature our team had scoped at four to six weeks. Cross-functional complexity, multiple stakeholder inputs, enterprise-grade quality requirements.

The difference wasn't the AI. We were already using AI. What most enterprise teams are still doing is what the methodology calls *vibe coding*: issuing prompts, generating output, iterating on instinct. It works for individual developers building something contained. At enterprise scale, with cross-functional teams, production-grade requirements and regulatory obligations, it produces inconsistent quality, governance debt, and a slow erosion of confidence in the code you're shipping — even as you ship it faster.

The difference with AI-DLC was structure. Specifically, what the methodology calls **collective context**. Before a line of code was written, the whole team — product, engineering, design — built a shared understanding of what we were actually trying to build and why. We gave the AI a complete picture rather than one developer's interpretation of a prompt. The output quality was categorically different. Not incrementally better. Different in kind.

The three phases that make it work

01 — INCEPTION

Humans build collective context

A deeply human phase. Cross-functional stakeholders create validated specifications and shared intent *before* a line of code is written. Quality of output downstream depends entirely on the thinking that precedes it.

02 — CONSTRUCTION

AI executes — humans oversee

AI works against validated specifications. Humans stay in the loop at the moments where quality, compliance and architectural integrity are at stake. The AI moves fast. Humans make the calls that matter.

03 — OPERATIONS

Continuous feedback loops

AI assists with deployment, monitoring and the feedback loops that keep production healthy. But what the signals mean, what to act on, and what warrants escalation remain human decisions.

The bottleneck was never the AI's capability. It was the absence of structure around it. **Give the AI collective context, keep humans in the loop at the right checkpoints, and the acceleration follows — without sacrificing quality or governance.**

PART 02 THE GOVERNANCE GAP

Nobody is talking about this yet.

Every product carries lifecycle obligations that don't expire at deployment. What changes at AI speed is not the obligation — it is the rate at which the gap accumulates.

PART 02 · THE GOVERNANCE GAP

What product lifecycle governance actually means

Product lifecycle governance is the ongoing discipline of understanding what you have built, what it is doing, what risk it carries, and whether it still belongs in your portfolio. It operates across the full arc of every asset's life — from inception through growth, maturity, change, and eventual sunsetting. It is not an event. It is a continuous practice that keeps organisations in command of their portfolio rather than subject to it.

Most organisations are held back not because they lack intent, but because the information they need is scattered across teams, functions and systems. Signals get lost in spreadsheets, email trails and disconnected workflows. Skyjed was built to change that — to move organisations from reactive to proactive, with a complete, shared view of performance and risk across every asset.

Four obligations attach to every product and asset, regardless of how quickly it was shipped:

Regulatory alignment	Does this asset comply with current and anticipated regulatory requirements? Is that compliance actively monitored, or assumed from the point of launch?
Strategic coherence	Does this asset still serve the strategy it was built to advance? Has the strategy changed in ways that make it redundant or misaligned?
Performance visibility	Is this asset performing as intended? Are the signals needed to detect performance changes early — and support proactive management — actually being monitored?
Sunsetting readiness	When this asset reaches the end of its useful life — or a better alternative exists — is there a clear, governed path to retirement? Or will it persist by default, accumulating cost and technical debt?

These obligations do not expire at deployment. They do not diminish because the development process was fast. **And they are not discharged by shipping confidently.**

FEATURE · THE FIVE FAILURE MODES

What accelerates when speed outruns governance.

The governance gap is the distance between what AI-driven development can ship and what governance infrastructure can oversee. Five failure modes accelerate without a corresponding investment in governance capability.

01 FAILURE MODE

Shadow product risk

Assets that exist in production but are not actively tracked, assessed or governed. Lifecycle drift sets in silently. At AI development pace, the shadow portfolio grows as fast as the team can ship.

02 FAILURE MODE

Decommission debt

Products and features that should be sunsetted are not — because no one has the visibility to identify them, the mandate to act, or the process to retire them safely. One of the most expensive, least visible liabilities in any portfolio.

03 FAILURE MODE

Regulatory lag

The development cycle used to create time for compliance review. AI-DLC eliminates that buffer. Assets with unreviewed regulatory exposure become a predictable output of AI-speed delivery.

04 FAILURE MODE

Ghost R&D

Effort invested in features that duplicate existing functionality, address solved problems, or serve strategic priorities that have since changed. Without lifecycle visibility across the full Asset Register, duplication compounds.

05 FAILURE MODE

Portfolio coherence breakdown

AI-speed development produces assets at a rate that individual product strategies — let alone enterprise portfolio strategies — were not designed to absorb. The result, over time: a portfolio where individual assets are well-built and the portfolio as a whole has lost strategic coherence.

> *The most significant governance failures in regulated industries rarely announce themselves. They accumulate — in the shadow portfolio, in unreviewed compliance exposure, in the sunsetting backlog no one has formally named. By the time the failure becomes visible, the liability is already substantial.*

PART 03 · THE FRAMEWORK PROBLEM

Existing governance was not built for this speed.

The product governance frameworks most enterprise organisations operate today were built incrementally, over years, in response to human-paced development. They are not reasonable now.

Stage gates were calibrated to weekly development cycles. Quarterly portfolio reviews assumed the portfolio changed slowly enough for a quarterly cadence to be meaningful. When the development cycle collapses from weeks to hours, the governance framework faces a structural choice: slow development to match governance capacity, or abandon governance to match velocity.

The instinct to cut governance is wrong

The instinct, when governance becomes the bottleneck, is to cut it. In regulated industries particularly, this is precisely wrong. Cutting governance does not eliminate the obligations it discharges — it delays the point at which they surface, and increases the cost of addressing them when they do.

What AI-speed governance actually requires

The answer is not to do governance more slowly. It is to build governance infrastructure that operates at the same speed as the development processes it oversees. Three capabilities define it:

01 **Continuous portfolio visibility**

A live Asset Register — real-time awareness of every asset, what lifecycle stage it occupies, what risk it carries, who is accountable.

02 **Automated health signals & smart reporting**

Health-check scores and course-correction triggers that detect when an asset is degrading, accumulating risk, or approaching sunsetting readiness — before those conditions become failures.

03 **Proactive lifecycle intelligence**

Analytical capability to assess full lifecycle cost, strategic alignment and portfolio coherence continuously — turning operational data into strategic decisions.

PART 04 · LIFECYCLE GOVERNANCE AT AI SPEED

Make the full lifecycle visible — from inception.

One of the most persistent governance failures in enterprise product management is assessing assets on the basis of their development and launch costs alone. The cost to build a feature is visible, attributed and reported. The cost to monitor it, maintain it, manage its regulatory obligations, and eventually retire it safely — these are frequently invisible, unattributed and unmanaged.

THE SKYJED FRAMEWORK

ABCd — Full-Cycle Product Economics

A	B	C	d
Acquire	**Build**	**Carry**	**decommission**
Investment to identify, scope and bring an asset to market.	Development and integration cost — where AI-DLC compresses delivery dramatically.	Ongoing performance monitoring, maintenance, regulatory compliance and oversight required to keep the asset healthy and protected from lifecycle drift.	Cost and risk of safe sunsetting — with full traceability and no residual exposure.

AI-DLC's primary impact is on **Acquire** and **Build**. The **Carry** obligations on a 48-hour feature are identical to a six-week one — and at AI development pace, assets are added faster than they are ever removed.

The Lifecycle Health Score

Skyjed's ML-driven Lifecycle Health Score provides continuous, quantified assessment of product and portfolio health. At AI velocity, a static, point-in-time assessment is insufficient — an asset aligned at launch can be misaligned, overdue for change, or accumulating regulatory exposure within months. The score is intelligence; what organisations do with it is judgment. **Act on what drives growth, not just what demands attention.**

Four dimensions of continuous health

- Strategic alignment to current portfolio direction
- Risk exposure across regulatory and operational vectors
- Performance against the metrics that matter
- Lifecycle stage appropriateness — is it where it should be?

PART 04 · SHADOW PRODUCT RISK

Shadow product risk at AI scale

Shadow product risk — the governance exposure carried by assets that exist in production outside active lifecycle oversight — is one of the most consequential and least discussed risks in enterprise portfolios. Every organisation has them: features built for use cases that no longer exist, products acquired and never fully integrated into the Asset Register, capabilities built by teams since reorganised with no clear current ownership.

> *At human development pace, the shadow portfolio grows slowly enough to be managed — imperfectly, belatedly, but managed. At AI development pace, it grows as fast as the team ships.*

Skyjed's shadow product risk framework makes the invisible portfolio visible: identifying assets outside active governance, quantifying the exposure they carry, and creating a structured path to bring them back under oversight — ensuring nothing falls through the cracks of the governed portfolio.

PART 05 · REGULATED INDUSTRIES

The regulatory reality

Regulators in financial services, insurance, and energy do not move at AI speed. They move at the speed of policy, consultation, and enforcement.

The EU Product Act, financial services conduct risk frameworks, and insurance product governance requirements share a common architecture: they place ongoing obligations on organisations to demonstrate active, documented governance of the assets they offer. The standard is not *"we reviewed this at launch"* — it is *"we maintain continuous governance of this asset's lifecycle, compliantly and at scale."*

AI-speed development intensifies that obligation: it increases the rate at which new assets enter the governed portfolio, and reduces the natural checkpoints at which regulatory alignment has historically been reviewed. The response is not to slow development down — it is to invest in continuous oversight, documented health-check scores and auditable lifecycle records with full traceability.

A leadership readiness checklist.

Seven questions for CEOs, CPOs and CTOs adopting AI-driven development. An honest measure of whether your lifecycle governance can keep pace with your delivery velocity.

01

Do you have a live Asset Register — real-time visibility into every product and asset in your portfolio?

Not a spreadsheet updated quarterly. A single, coherent view of what exists from inception to sunsetting, who owns it, what lifecycle stage it occupies, what it costs to carry.

YES PARTIAL NO

02

Can you identify which assets carry active regulatory exposure today?

Not at launch — today. As regulations evolve and usage patterns change, the regulatory profile of an asset changes. Governance that cannot track that change in real time is not governance at AI speed.

YES PARTIAL NO

03

What is the full ABCd lifecycle cost of what you shipped in the last 90 days?

Development cost is visible. Carry cost — performance monitoring, compliance overhead, operational integration, eventual sunsetting — frequently is not. Can you quantify it?

YES PARTIAL NO

04

How many assets have not been formally reviewed in the last 12 months?

For most organisations, the answer is more than they expect. Lifecycle drift is real. Quantifying the shadow portfolio is the first step to governing it.

YES PARTIAL NO

05

Do you know which assets are candidates for sunsetting — and what the cost of not sunsetting them is?

Decommission debt is one of the most expensive and least visible liabilities in a portfolio. Is it actively managed, or accumulating by default?

YES PARTIAL NO

06

If your regulator asked for a complete product risk assessment today, how long would it take?

Days is concerning. Weeks is a governance failure. The answer is a direct measure of governance infrastructure maturity.

YES PARTIAL NO

07

Is your product portfolio more or less strategically coherent than 12 months ago?

AI-speed development adds capability rapidly. Without active coherence assessment, drift may not become visible until it creates a strategic or regulatory problem.

YES PARTIAL NO

If you can answer all seven with current data, your governance infrastructure is positioned for AI-speed development. For most organisations, the exercise of attempting to answer them is itself a governance gap assessment.

CONCLUSION

The next competitive advantage is governed speed.

I want to close with something I genuinely believe, having built a platform on both sides of this equation.

The organisations that build the fastest in the next five years will have a significant advantage. That is not in question. What is in question is whether speed alone is sufficient.

Growth has never been linear. Companies don't grow in phases — they grow through feedback loops: the recurring cycle of signals, issues, reviews and opportunities that shape what needs to happen next. The organisations that win are the ones that can see change early, accelerate those feedback loops, and act with confidence — not the ones that simply execute the fastest.

AI-driven development has changed the development half of that equation fundamentally. The 15–20× acceleration is real, validated and spreading. The governance half has not kept pace — and for most organisations, it will not unless governance is treated as a strategic investment rather than a compliance cost.

> *Speed and governance are not opposing forces. The discipline that makes AI-DLC work — collective context, validated specifications, human oversight at the moments that matter — is the same discipline that makes product lifecycle governance work. Both depend on humans being in the right place in the process.*
>
> — STEPHEN BROWN, CO-FOUNDER & CTO, SKYJED

You can now build at 15–20× speed. The question is whether your organisation is equipped to govern what you build — at the speed you build it.

THE NEXT CONVERSATION

See what governed speed looks like on your portfolio.

Skyjed gives leaders complete visibility and oversight over the full lifecycle of every product or asset — at the velocity AI-driven development now demands.

BOOK A DEMO

skyjed.com/demo

Scan the QR code or visit the link to book a 30-minute conversation with the Skyjed team.

ABOUT SKYJED

Skyjed is an AI-powered product lifecycle governance platform built for enterprises in regulated industries — including financial services, insurance, and energy. Founded in 2019 following more than 100 global CEO interviews validating the product lifecycle governance gap, Skyjed provides the continuous portfolio visibility, automated health signals, and lifecycle intelligence enterprise organisations need to govern at AI speed.

Skyjed's proprietary frameworks — including the Lifecycle Health Score, Full-Cycle Product Economics (ABCd), and shadow product risk methodology — are purpose-built for the complexity of enterprise product portfolios operating under regulatory obligation and competitive pressure.

Visit
skyjed.com
skyjed.com/demo

Connect
Stephen Brown on LinkedIn
Released at AWS Summit Sydney 2026

Regions
Australia · United Kingdom · Europe