

Vibe Coding and Enterprise Security

Why we're all-in, and what
we built to support it safely.

Vibe Coding and Enterprise Security


The shift is real. Vibe coding (the practice of using AI to generate and ship production-grade code at speed) has moved from experiment to mainstream faster than most enterprise technology teams anticipated. Developers are going from brief to deployment overnight. Whole features are being scaffolded, tested, and pushed to production in timeframes that would have been implausible eighteen months ago.

At Dataweavers, we think this changes the economics of software delivery in ways that are genuinely positive, for the right teams, on the right foundations.

Speed amplifies everything underneath it

When AI assists in generating production code at velocity, it doesn't discriminate between the strong parts of your codebase and the weak ones. Every pattern gets amplified. The solid architecture, the clear deployment contracts, the observable pipelines, those scale brilliantly. But so do the gaps. Security assumptions that were never quite right. Observability that was bolted on rather than built in. Deployment pipelines that were designed for human-paced release cycles, not overnight AI-assisted shipping.

Speed without a strong foundation doesn't just expose problems. It accelerates them. The failure modes that used to surface gradually now arrive at scale, and often faster than a team can respond.



**The teams getting the most from
vibe coding are the ones who sorted
their platform operations first.
The foundation is the multiplier.**



What agentic AI actually asks of your infrastructure

The AI features inside modern DXPs are no longer conceptual. Agentic publishing workflows, AI-assisted personalization at scale, dynamic content orchestration driven by real-time signals, these are shipping now, and the roadmap accelerates from here. In fact, [McKinsey survey respondents](#) expect AI agent adoption to double in the next three years, with nearly 50% expecting AI to be embedded across the full technology stack. The infrastructure question is whether your platform was designed for that reality.

But the AI layer doesn't sit on top of your architecture. It sits inside it, interacting with your rendering layer, your content APIs, and your deployment pipelines in ways that a platform assembled for a 2023 headless launch wasn't designed to handle.

Most enterprise stacks have one of the three things they need. [Gartner predicts over 40% of agentic AI projects will be cancelled by end of 2027](#), not because the AI isn't capable, but because of inadequate risk controls and unclear operational foundations. The platform underneath is where these projects succeed or fail. Here's what genuine readiness looks like.

1. Real-time observability across every layer

Uptime monitoring is not observability. What agentic workflows require is deep, continuous visibility into your rendering stack, your APIs, your content pipelines, at the moment things are happening, not in a dashboard you check the next morning. Agentic AI generates failure modes that weren't in anyone's incident runbook twelve months ago. The teams that catch them early are the ones whose infrastructure was instrumented to surface them in real time, not discovered by a customer complaint.

What can you do? Map your blind spots today. List every layer your rendering stack touches from APIs to pipelines, CDN to content delivery, and ask honestly: if something failed in each one right now, how long before you'd know? Anything longer than minutes is a gap worth addressing before an agentic workflow finds it first.



2. Compliance-ready pipelines, built in from the start

When AI agents begin touching content workflows: drafting, publishing, personalizing, deploying, the compliance surface area expands significantly. Every action needs a trail. Every deployment needs to be auditable. Regulated industries and enterprise procurement processes increasingly require this evidence, and the expectations are only moving in one direction.

Retrofitting audit capability onto a platform that wasn't designed for it is a significant undertaking. The teams in the strongest position are the ones whose compliance scaffolding was architectural, not added after the fact under pressure.

What can you do? Run a shadow audit. Pick one content workflow and trace every action an AI agent would take through it - draft, approve, publish, deploy. Ask whether each step produces a timestamped, attributable record. If you can't answer that question cleanly, you don't have compliance infrastructure, you have compliance intention.

3. A platform designed to be operated, not just launched

There is a meaningful difference between a headless platform that was assembled to reach go-live and one that was designed to be run in production over time. The former optimizes for the launch milestone. The latter optimizes for what comes after, the upgrade cycles, the incident response, the scaling events, the security patches, the operational ownership that accumulates as a platform matures.

Agentic AI asks far more of the infrastructure underneath it than traditional DXP usage patterns. Stability, predictability, and clear operational ownership aren't optional features. They're the prerequisites that make everything else possible.


What can you do? Ask who owns 2am. For every critical layer of your platform, name the person or team who is accountable when something breaks outside business hours. If the answer is unclear, or if it's "the vendor," that's your operational ownership gap, and it's the one agentic AI will expose fastest.



What Dataweavers built for this moment

We have spent years building a managed platform layer that sits between the DXP and production, one that was designed around operability from the beginning, not retrofitted toward it. The observability, the compliance scaffolding, the deployment discipline, the security posture: these aren't add-ons we sell. They're the architecture.

That means when our clients move into AI-native DXP territory, adopting agentic workflows, enabling AI-assisted publishing, scaling personalization, they're arriving at it with a stable, observable, compliant platform already running underneath. Not because AI is a risk to be managed. Because a solid foundation is what lets you move fast on the things that actually matter.



Vibe coding is a genuine shift.
The infrastructure question is
whether your platform was designed
to support what comes with it.

Take AI-Readiness to AI-Reality

If your team is evaluating platform readiness ahead of an AI-native DXP rollout, we'd like to talk. We can walk you through what enterprise-grade operability looks like in practice, and where the gaps most commonly appear in stacks that weren't built for this moment.

Fix Less. Build More. | hello@dataweavers.com



About Dataweavers

Dataweavers is a Microsoft-first, Azure-native platform company that helps enterprise organisations run complex digital platforms at scale.

From front-end hosting for Next.js application delivery to infrastructure, security, and DevOps, we handle the platform operations that enterprise teams shouldn't have to manage themselves.

Backed by 24x7 global support, we give enterprise teams the freedom to fix less and build more.

