

Reimagining the Software Development Lifecycle in the Age of AI

by

Jason Brucker
Managing Director, CIO Solutions

Scott Laliberte
Managing Director, CISO Solutions

protiviti®
Global Business Consulting

Artificial intelligence (AI) is reshaping software delivery, though its most significant effects extend well beyond coding efficiency. Early adoption often focused on accelerating discrete tasks — especially coding. Those gains are real, but they represent only a fraction of AI’s impact. High-performing organisations are recognising that AI changes how work flows across the entire software development lifecycle (SDLC), from idea intake through value realisation.

AI does not replace agile or modern engineering disciplines. Instead, it exposes them. In practice, AI behaves as an amplifier — magnifying the strengths of mature delivery systems and the dysfunctions of struggling ones. When throughput increases without corresponding improvements in quality, governance and risk controls, the organisation experiences higher rework, instability and security exposure.

Protiviti’s [AI Pulse](#) research underscores an additional reality: Organisations that embed AI into workflows and operating processes — supported by governance and measurement — are more likely to exceed ROI expectations. Data confidence is a major differentiator; organisations confident in their data are significantly more likely to exceed their AI ROI targets.

What Has Changed

- AI shifts work from implementation effort to decision quality (architecture, validation, governance and accountability).
- AI-enabled products are rarely “done” — models drift, data shifts and user behaviour changes, requiring continuous learning and adaptation.
- AI expands the risk surface (data exposure in prompts and outputs, probabilistic behaviour, new supply chain dependencies), raising the bar for governance by design.
- Metrics must evolve from local team output to end-to-end value streams (cycle time, stability, AI rework and cost per feature).

RESEARCH INSIGHT

“AI’s primary role in software development is that of an amplifier. It magnifies the strengths of high-performing organizations and the dysfunctions of struggling ones.”

– [DORA, State of AI-Assisted Software Development 2025](#)

Why AI Changes the SDLC (Beyond Coding Productivity)

AI reshapes the SDLC because it changes decision cycles, increases delivery velocity and makes quality issues propagate faster. In traditional delivery systems, delays and friction often hide weaknesses — unclear requirements, inconsistent testing, poor traceability, brittle architectures and late-stage security checks. AI reduces friction in the build phase, which can expose bottlenecks and control gaps elsewhere in the lifecycle.

AI does not create delivery maturity; it rewards it. When governance and quality systems are strong, AI increases leverage. When they are weak, AI increases the speed of failure.


Agile still matters — now it must function as a learning system

Agile principles — incremental delivery, rapid feedback and adaptability — remain foundational. What changes is how those principles must be applied when AI contributes meaningfully to design, development and validation. What's more, models evolve, data distributions shift and user behaviour changes continuously, creating an ongoing cycle of refinement rather than a finite project end state.

This places renewed importance on agile ceremonies — not as process formalities but as learning engines. Among these, the retrospective becomes critical. Retrospectives become structured reflection points where teams examine how AI behaved in the prior iteration: where automation improved flow, where AI introduced rework or instability, and how prompts, guardrails, validation steps or workflows should evolve.

From project-centric to product-centric delivery

AI accelerates the shift from project-centric delivery toward product-centric delivery. Product management increasingly replaces traditional project management within the SDLC, focusing on continuous value realisation rather than fixed milestones. At the same time, enterprise program management remains essential for large strategic initiatives, providing investment governance,



dependency management and executive oversight in alignment with AI-enabled delivery.

Building AI-Ready Delivery Organisations

As organisations move beyond early AI experimentation, many discover that localised productivity gains do not automatically produce enterprise-level outcomes. The constraint is rarely technology. It is organisational design — roles, workflows, governance and learning mechanisms — that determine whether AI becomes scalable advantage or isolated efficiency.

Role convergence and the shift to decision-centric work

Protiviti's [AI Pulse](#) research highlights a progression: As AI maturity increases, organisations shift focus from task-level efficiency to decision-making, governance and human oversight. As AI assumes more implementation work, roles converge and skill demands shift. Developers focus less on writing boilerplate code and more on validation, architecture and judgment. Quality engineering moves toward strategy and automated quality systems. Security, risk and audit functions move earlier into delivery. Product owners increasingly anchor prioritisation and continuous value realisation across cross-functional teams.

SDLC discipline becomes more important — not less

AI changes how work is performed, but it does not remove the need for disciplined SDLC practices. Clear requirements, traceability, testing rigor, separation of duties and controls remain essential — both for building AI-enabled systems and for using AI-assisted tooling safely within the delivery lifecycle.

- **Requirements and traceability:** Make scope, acceptance criteria and nonfunctional requirements explicit; link changes to value streams and controls.
- **Testing rigor:** Expand automated test coverage (unit, integration, security, performance) to absorb higher change volume.

AI PULSE HIGHLIGHT

“Organizations confident in their data are 3x more likely to exceed AI ROI expectations.”

– [Protiviti AI Pulse Survey 2025 \(Vol. 2\)](#)

- **Separation of duties:** Preserve independent review and approvals — especially for high-risk changes or regulated systems.
- **Auditability:** Maintain evidence (what AI produced, what humans approved and how validation was performed).

Governance by design: Enabling speed without inviting shadow AI


Governance embedded directly into delivery workflows enables speed rather than constraining it. One-size-fits-all AI controls often backfire because AI systems touch sensitive data (prompts, outputs, training and tuning), behave probabilistically, depend on complex supply chains and evolve rapidly. Heavy reviews for every AI experiment slow learning and invite work-arounds; permissive approaches increase exposure.

Protiviti's work with CISOs points to a pragmatic answer: Apply streamlined fast paths for low-risk AI use cases and pre-vetted tools, while requiring deeper diligence for moderate- and high-risk scenarios. This creates a sanctioned, repeatable path so teams do not route around security and governance.

Risk signals reinforce why governance must scale with adoption. Protiviti notes that the number of AI projects in production has increased rapidly year over year, and IT audit leaders are increasingly treating AI as a significant technology risk over the next two to three years.

Governance by Design: AI-Accelerated Security Risk

Recent advances in AI-driven offensive security raise the bar for how organisations govern the SDLC. Capabilities such as [Anthropic's Claude-based Mythos](#), a recent application of large language models to security research, show how quickly AI can identify, link, and exploit vulnerabilities across applications, code and infrastructure. As these tools mature, the time between [vulnerability discovery](#) and exploitation continues to shrink, exposing the limits of manual reviews and late-stage security checks.



For leaders, the implication is clear: Security must move at the same pace as delivery. Organisations are embedding AI-assisted security analysis directly into development pipelines, using models such as Claude to surface risk earlier, before code reaches production. When implemented with the right controls, these reviews run continuously, add little friction and improve coverage. Using AI in the SDLC to counter AI-enabled attacks is no longer a differentiator. It is becoming the minimum standard for maintaining resilience as security shifts to machine speed.

Measuring What Matters: KPIs for an AI-Enhanced SDLC

AI changes how value is created in software delivery, and metrics must change with it. Traditional SDLC measures were designed to optimise local efficiency (velocity, throughput and sprint completion). In AI-enabled environments, these measures no longer explain enterprise impact, risk exposure or return on investment. AI accelerates delivery while increasing the cost of failure, making end-to-end visibility essential.

In AI-enabled delivery, even small units of work — such as a single enhancement or defect fix — must be understood in the context of an enterprise value stream. Value streams represent the end-to-end flow of work from idea to customer or business outcome. When activities are explicitly connected to value streams, leaders can see where AI is creating leverage — or introducing risk.

RESEARCH INSIGHT

“The greatest returns on AI investment come not from the tools themselves, but from a strategic focus on the underlying organizational system.”

– *DORA, State of AI-Assisted Software Development 2025*

Core metrics for an AI-enhanced SDLC (executive view)

AI increases throughput and the cost of failure. The following KPI set helps leaders track end-to-end flow, stability, quality, AI effectiveness and enterprise value.

Metric Category	Metric	How It's Calculated	Why It Matters
Flow efficiency	Idea-to-production cycle time	Time from approved concept to production	Measures enterprise delivery speed
Stability	Change failure rate	Failed releases ÷ total releases	Ensures that speed does not increase risk
Quality	Postrelease defect rate	Defects in production ÷ releases	Protects customer and brand impact
AI effectiveness	% AI-generated code reworked	Reworked AI output ÷ total AI output	Indicates AI quality and prompt maturity
Human oversight	AI suggestion acceptance rate	Accepted ÷ reviewed suggestions	Measures trust and validation effectiveness
Enterprise value	Cost per feature	Delivery cost ÷ features delivered	Connects delivery to ROI

Use metrics as learning inputs — not performance theater

In high-performing organisations, metrics serve as learning inputs rather than static performance reports. Executive dashboards balance speed, quality, stability and value. Retrospectives at the team and program levels use these metrics to guide continuous improvement — especially around AI rework, workflow friction and risk hot spots.

A practical SDLC blueprint for AI-enabled delivery

C-suite leaders can accelerate maturity by standardising an AI-aware SDLC blueprint. The objective is not to add bureaucracy; it is to make speed safe, repeatable and auditable.

- **Idea intake and triage:** Classify AI use cases by risk and data sensitivity; require a lightweight intake form and an inventory of AI-enabled components.
- **Design and architecture:** Define data boundaries, model dependencies and control requirements early; design for monitoring and drift management.
- **Build:** Use AI assistance where it improves flow, but require peer review and automated checks; capture evidence of AI-generated artifacts.
- **Test and validate:** Expand automated testing; validate AI outputs for correctness, security and (where applicable) bias.
- **Release:** Maintain separation of duties; ensure that release evidence is complete; use progressive delivery where appropriate.
- **Operate and improve:** Monitor production behaviour, drift and anomalies; feed learnings back into retrospectives and backlog prioritisation.

PRACTICAL GOVERNANCE MODEL

“Use a simple Green/Yellow/Red zoning approach: fast-path low-risk use cases, require due diligence for moderate risk, and apply high scrutiny (human-in-the-loop, auditability and protected environments) for consequential decisions.”

– *Protiviti, Pragmatic AI Security Strategies for CISOs*

What Executive Leaders Should Do Next (A 90-Day Plan)

The fastest path to sustainable AI-enabled delivery is to treat SDLC modernisation as an operating-model program, not a tooling rollout. The following 90-day actions are designed to produce measurable progress without disrupting delivery.

For CIOs, CDOs and CTOs

- **Map and instrument priority value streams** from idea intake to production and value realisation; ensure that AI-accelerated work is visible end to end.
- **Strengthen the learning system:** Enforce retrospectives that explicitly assess AI impact (flow, rework, stability) and capture changes to prompts and guardrails.
- **Rebalance roles toward judgment:** Validation, architecture, quality strategy and accountability must scale with AI throughput.
- **Standardise an AI-aware SDLC** blueprint and publish golden paths for teams (platform tooling, test automation, observability and release patterns).

For CISOs

- **Operationalise risk-based zoning** (Green/Yellow/Red) with intake forms, inventories and one-page playbooks that define required controls and turnaround times.
- **Require protected environments** for sensitive data (tenant isolation, encryption, access controls, DLP, logging and redaction).
- **Mandate human-in-the-loop controls** for high-impact decisions and define accountability, override and recourse mechanisms.
- **Implement continuous monitoring** and periodic red-teaming for AI systems that affect customers, finances, safety or regulated decisions.

- **Ensure that AI-based security review** is built into the SDLC process.

For the broader executive team

- **Replace local productivity reporting** with an AI-enhanced SDLC dashboard that balances speed, stability, quality, AI rework and enterprise value.
- **Align governance expectations to risk appetite:** Move quickly where risk is low and invest diligence where impact is high.
- **Treat data confidence as a foundation** for trustworthy AI and ROI (governance, training, transparency and quality at data entry points).

Conclusion

AI will continue to improve. The differentiator will not be which model you choose; it will be whether your SDLC — operating model, governance and measurement — can convert AI capability into enterprise value safely and repeatedly. Organisations that combine disciplined delivery, pragmatic governance and outcome-oriented metrics can turn AI-driven speed into sustained advantage — without sacrificing trust, quality or control.

References

- “AI and the Future of Software Engineering: Faster Development, Higher Risk, New Accountability,” *The Protiviti View*, March 13, 2026, <https://blog.protiviti.com/2026/03/13/ai-and-the-future-of-software-engineering-faster-development-higher-risk-new-accountability>.
- *AI Pulse Survey 2025 Vol. 1: From AI Exploration to Transformation*, Protiviti, www.protiviti.com/us-en/survey/ai-pulse-vol1.
- *AI Pulse Survey 2025 Vol. 2: From Data Confusion to AI Confidence*, Protiviti, www.protiviti.com/us-en/survey/ai-pulse-vol2.
- *Pragmatic AI Security Strategies for CISOs*, Insights paper, Protiviti, www.protiviti.com/sites/default/files/2026-04/pragmatic-ai-security-strategies-for-cisos_e.pdf.
- *State of AI-Assisted Software Development 2025*, DORA, <https://dora.dev/research/2025/dora-report>.

About the authors



Jason Brucker
Managing Director, CIO Solutions

Jason is a seasoned technology strategist with over 25 years of experience and has worked with a broad range of global clients to help them achieve their enterprise transformation goals. As an adviser to senior technology leaders, Jason and his teams have successfully delivered targeted assessments and advisory expertise as well as multiyear system and operational transformation initiatives.



Scott Laliberte
Managing Director, CISO Solutions

Scott enables clients to leverage emerging technologies and methodologies to innovate, while helping organisations transform and succeed by focusing on business value and managing risk. His team specialises in many technological areas, including artificial intelligence (AI) and machine learning, the Internet of Things (IoT), the cloud, blockchain and quantum computing.

Protiviti (www.protiviti.com) is a global consulting firm that helps clients transform and protect their businesses, and respond to planned and unexpected events. Through a network of more than 90 offices in over 25 countries, Protiviti and its independent and locally owned member firms deliver deep expertise and tailored capabilities across technology, artificial intelligence, data, operations, finance, legal, compliance, HR, marketing, digital, risk, and internal audit—enabling organisations to accelerate innovation, navigate risks and safeguard what matters most.

Named to the [Fortune 100 Best Companies to Work For](#)® list since 2015, Protiviti Inc. has served more than 80 percent of Fortune 100 and nearly 80 percent of Fortune 500 companies. The firm also works with government agencies and smaller, growing companies, including those looking to go public. Protiviti Inc. is a wholly owned subsidiary of Robert Half (NYSE: RHI).