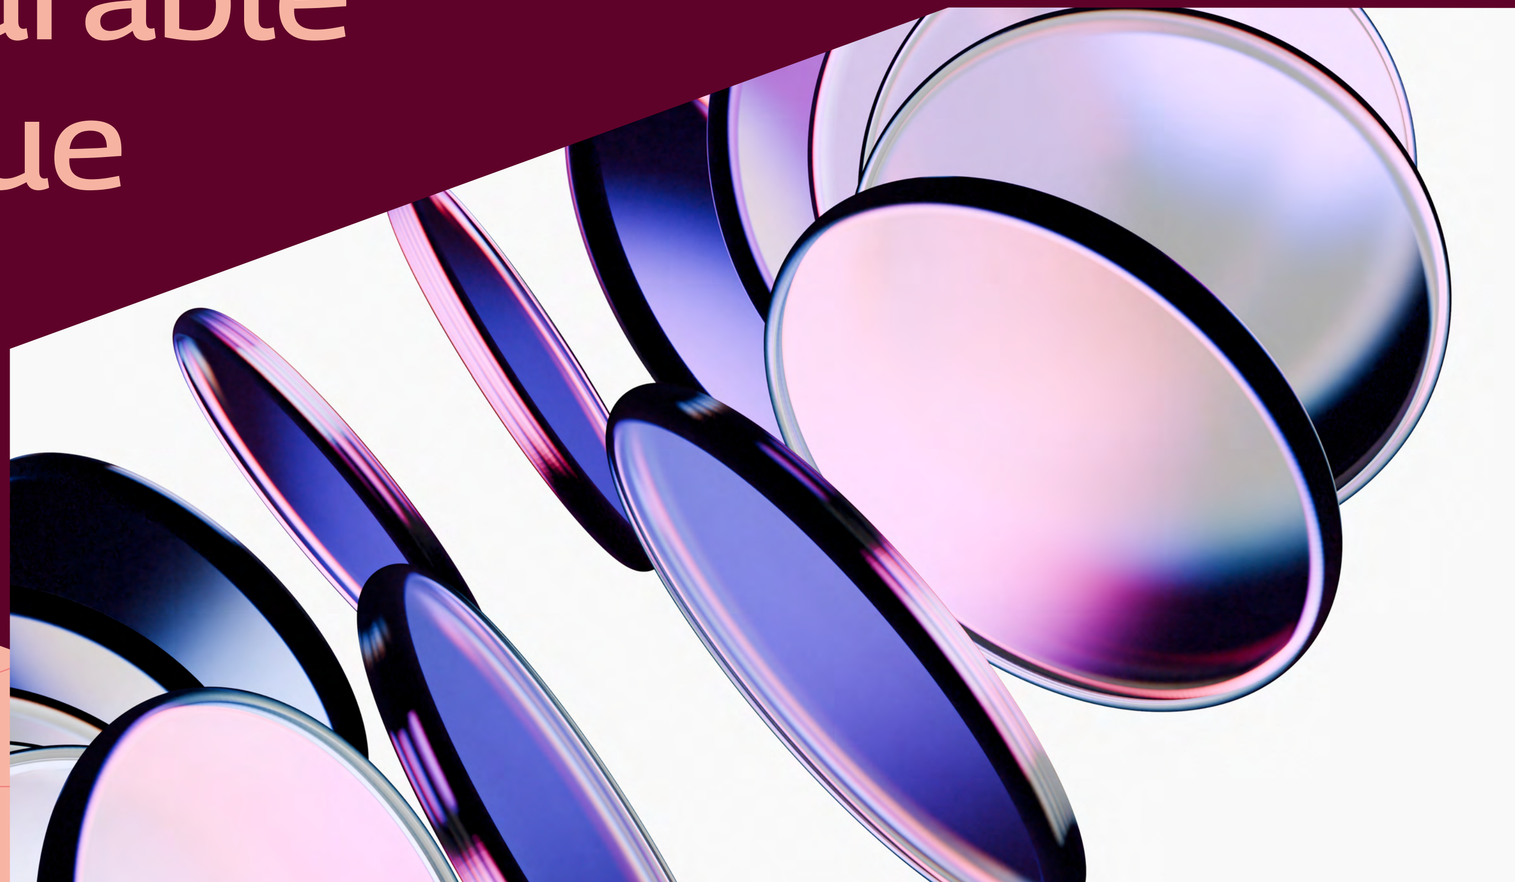
Whitepaper Read Time - 8 Mins.

TECH mahindra

Agentic AI SDLC: ROI and Measurable Enterprise Value



Executive Summary

Agentic AI is playing a large role in transforming the economics of software development, with AI agents increasingly taking point on key tasks across development, architecture, testing, release, and operations, automating tasks that previously required significant human effort, not just in code generation but also in analyzing dependency gaps and enabling modernization.[1]

First, let us make the important distinction that productivity improvement is not the same thing as enterprise value. And the latter is what matters more to organizations investing in agentic AI. The most important question, as always, is whether improvements across the software development lifecycle (SDLC) can be measured and further amplified.

Realizing agentic AI ROI requires much more than just deploying AI development tools. It needs organizations to understand where true value is created across the SDLC and to have an acute knowledge of agentic capabilities, gaps, and points where value is lost. More importantly, it requires an understanding of how governance can help balance productivity gains with legal and ethical considerations.

This whitepaper examines that value chain and outlines a measurement-led path from agentic AI adoption to credible ROI.



Table of Contents

Title	Page
From Software Development to an Agentic AI Value Chain	4
Where Agentic AI Creates Value Across the SDLC	5
Finding the Right Agent for the Right Task	6
How SDLC Gains Need to Scale with Agentic AI	7
How Governance Is a Critical Element	8
A 90-Day Path to Measurable Enterprise Value	9
Five Imperatives for Measurable Agentic AI Value	10
Conclusion	11

From Software Development to an Agentic AI Value Chain

Much current discourse around AI agents has been primarily focused on coding productivity gains, but that is only a part of the equation. Overall, software development is an extensive, connected lifecycle, with every phase (requirements, architecture, development, testing) intertwined.[2]

Agentic AI therefore needs to be evaluated across the entire SDLC value chain.

Design → Development → Test and QA → Release

Every stage of the SDLC value chain presents an opportunity to rethink how ROI can be delivered through agentic AI, in terms of both economic impact and enterprise value.



Where Agentic AI Creates Value Across the SDLC

AI agents do not just enhance a single productivity layer; they enhance several. Depending on the nature of tasks being performed, the economic potential, and enterprise value vary significantly. The results can be seen across the entire SDLC value chain:



Design

Value is created in design by reducing the effort required from business analysts, solution architects, and senior engineers. It is, after all, a relatively small portion of the total FTE pool in the source model, but it involves higher-cost roles.

However, these roles still require strong human oversight. An inaccurate requirement, or poorly reviewed design decision can create downstream rework that offsets the initial productivity gain.



Development

Agentic AI can help reduce time and effort in development, which accounts for 60% of the total FTE pool and is the largest value pool. However, agentic AI value depends not just on adoption but on assigning the right capability to the right workload.



Test and QA

Higher-quality testing means reduced rework and defects, which translates to less time spent and eventually higher ROI. No matter the nature of agentic AI, the goal is to ensure that quality remains high in all cases.

Quality benchmarks on measured test-suite generation tasks can differ across agents. The objective is to keep them high, since it is an important part of the ROI equation.



Release

Release is the final stage of the value chain because it connects productivity with operational outcomes. Significant cost savings can be achieved when incident volumes are lower and recovery times are faster, especially in complex environments where identifying risk and errors is far more challenging and time-consuming.

Finding the Right Agent for the Right Task

The myth that enterprises can adopt a single AI development platform for their entire operations does not hold up in practice. The best path is to adopt a hybrid model across the enterprise.

Different AI agents suit different enterprise needs, and each performs best on specific tasks. The goal is always to deploy them to the tasks across the SDLC they are best suited to.



Here are some of the most prominent platforms and their best-use scenarios:

Platform	Best Suited For	Key Consideration
GitHub Copilot Enterprise	Daily workflows, inline completion, boilerplate, PR automation, broad IDEs	Strong fit for everyday development at scale
Claude Code	Complex refactoring, cross-service debugging, migration, architectural	Greater depth for context-heavy, complex workloads
Windsurf	Single-environment iteration and repository-aware execution	Maintain current deployment rather than expand (pending acquisition/assurance)

Figure 1: Internal Analysis of Agentic AI Platform Performance Across the SDLC

A strategic focus on economic alignment between workload, capability, cost, and outcome is what will determine the best hybrid strategy for deploying agentic AI across the SDLC. Organizations need to understand how to adapt for optimal impact and scale it as logically as possible.

How SDLC Gains Need to Scale with Agentic AI

Once value has been demonstrated at the task and program levels, the next challenge is scaling.

Often, potential value can be created across the SDLC but still fail to reach the bottom line, especially when organizations begin to rapidly expand the scale of their agentic AI implementation. Once that happens, they also increase review demand, pull requests, and governance requirements. Eventually, they end up acquiring a lot of extra technical debt.[3]

The objective, therefore, is not simply to maximize AI-generated output. It is to maximize valuable, reliable output. It is important to ensure that as adoption levels increase, so does the estimated net annual value in direct proportion. Successful adoption always skews positive and upwards.

Here are three enterprise scenarios observed during adoption across a population of 1,000 developers internally at Tech Mahindra.

Adoption Level	Estimated Net Annual Value
20% Adoption	$17.0M
50% Adoption	$27.3M
75% Adoption	$37.2M
Mid-Scenario Planning Range	$27.3M-$31.1M

Figure 2: Adoption Level of AI Agents Corresponding to Estimated Net Annual Value Increase

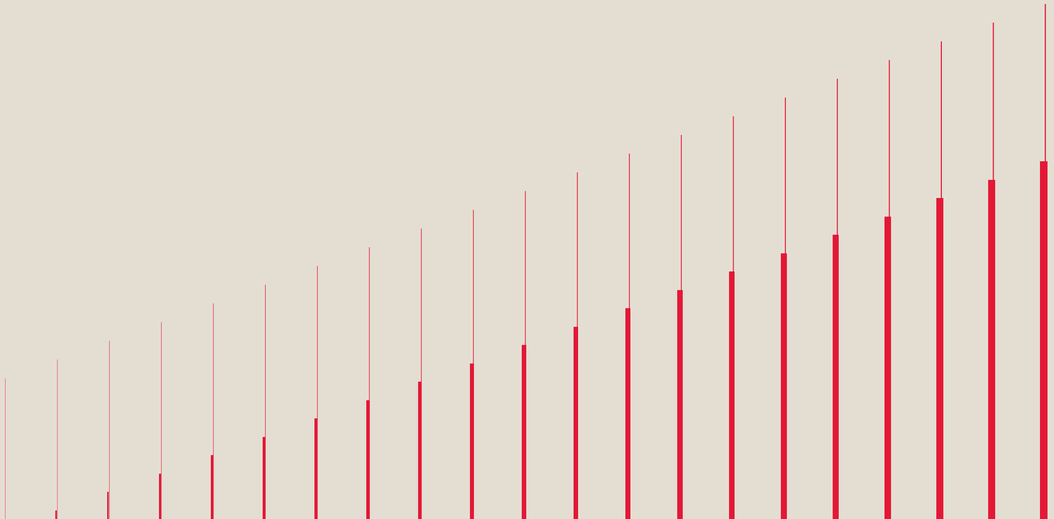


How Governance is a Critical Element

Now we come to governance, which is, simply put, the mechanism that protects the value of the SDLC rather than limiting it, as is sometimes assumed.

In fact, it offers the critical control layer for organizations as they deploy AI agents across the entire spectrum of the SDLC. It makes sure that productivity gains do not come at the cost of quality, security, ethical compliance, and accountability. It establishes clear policies, measurable quality thresholds, review gates, logging, and human accountability around AI-assisted work.

The Tech Mahindra SDLC governance framework optimizes costs without significant quality loss.



Governance Approach	What It Governs
Cost Takeout	It includes hard-dollar outcomes like cycle-time reduction, productivity, eliminated rework, and prevented incidents.
Code Quality	It tracks measures including defect escape, test coverage, PR review time, and vulnerabilities.
Human-in-the-Loop Oversight	It provides policy, logging, and review gates around AI-generated work.

Figure 3: The Tech Mahindra Governance Approach for Agentic AI Across the SDLC

Both the “cost takeout” and “code quality” approaches are enveloped by a “human-in-the-loop oversight” model. This provides a check on all metrics and data regarding outcomes, defects, vulnerabilities, and other relevant data points.

The goal, as always, is continuous learning and improvement. This involves learning from what works, what can be adapted to scale, what does not work, and how things can be improved for more quality control.[4]

A 90-Day Path to Measurable Enterprise Value

Here is how a human-led framework for agentic AI in SDLC comes into play. The Tech Mahindra SDLC playbook proposes a 90-day path that starts with governance, moves to instrumentation, and emphasizes using results to choose further actions.

Days 1–30: First, the goal is to clearly define governance and establish baselines. This involves identifying how AI assistance works as a whole and how its overall impact can be measured quantitatively and qualitatively.

Days 31–60: Now it is time to move to instrumentation, which involves tracking how AI is performing across the SDLC. This can be measured through dashboards, tool telemetry, monitoring, and DevOps Research and Assessment (DORA) baseline metrics to give a clearer overall picture.

Days 61–90: At this final phase, the organization should now possess enough evidence to run comparative readouts and compare AI-assisted and non-AI-assisted performance across the spectrum.

Timeline	Focus	Key Actions
Days 1–30	Establish Baselines	Define governance, identify AI-assisted work, and establish mechanisms to measure its impact accurately.
Days 31–60	Build Instrumentation	Implement dashboards, tool telemetry, quality metrics, review-time monitoring, and DORA baselines to track AI performance across the SDLC (Software Development Life Cycle).
Days 61–90	Comparative Readout	Compare AI-assisted and non-AI performance across tooling cost, quality, governance compliance, and developer experience.

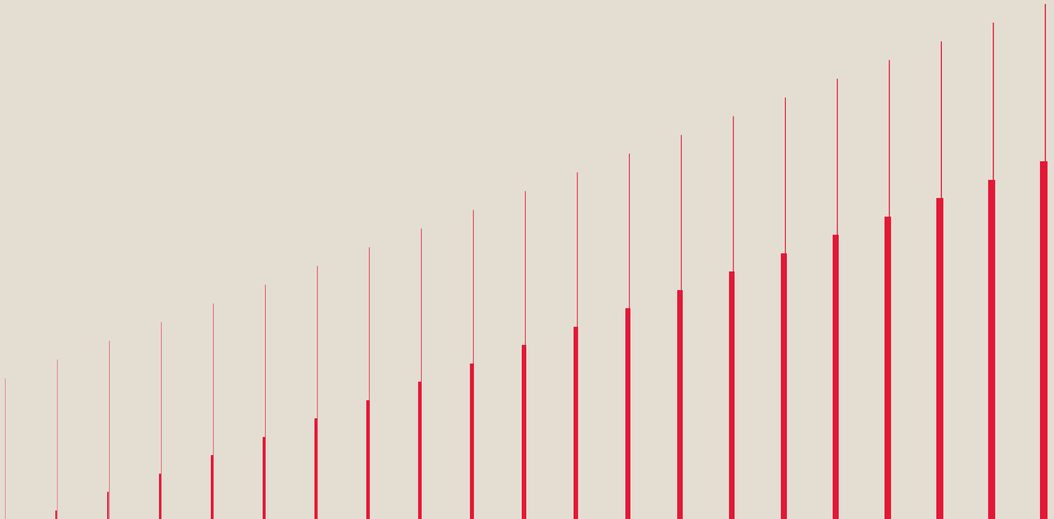
Figure 4: The 90-Day Path to Measurable Enterprise Value for Agentic AI in the SDLC

Based on the decisions being made beyond these 90 days, decision-makers and key stakeholders can propose whether to implement a broad transition or to choose selective deployment. All of this must be done logically and be evidence-based to successfully scale operations through agentic AI.[5]

Five Imperatives for Measurable Agentic AI Value

Moving from agentic AI experimentation to measurable enterprise value ultimately comes down to a few key disciplines. All of these are critical to address before implementation:

1. **Measure across the full SDLC:** It is important to measure ROI across all phases, not just coding but architecture, testing, release, and operations.
2. **Use economics appropriate to the phase:** Cost and expenditure have to be allocated in accordance with the value that is being created, or will be created.
3. **Human oversight has to be key to governance:** Build a “human-in-the-loop” oversight to monitor the agentic AI operations.
4. **Distinguish AI-assisted from non-AI-assisted work from Day 1:** From the outset, this makes it easy to build a comparative layer.
5. **Create clear targets for Day 90:** By Day 90, organizations should have enough critical data to determine how to best prepare for future deployment.





Conclusion

Deploying agentic AI across the SDLC is a thoughtful, strategic, and scalable practice, not just a one-time commitment. It is not only about deployment, but learning how to best measure, track, and sustain it at scale.

The evidence from the Tech Mahindra delivery program, per internal records, indicates a strong outlook. But the bottlenecks and constraints that remain must be recognized and organizations must stay current with governance rules and potential overheads.



Tech Mahindra Centric: Key Value Indicators

Value Metric	Measured / Modeled Impact
Annualized Savings	$2.45M
Planned FTEs Saved	14 of 24
Production Programs	3
Directional Year-1 Enterprise Value	$13.5M - $19M
Directional Steady-State Value at 50% Adoption	$27.3M - $31.1M

Once all that is considered, organizations will have a strong foundation to measure their ROI accurately and scale with confidence.



Srinivas Burra

Vice President and Head of
North America Consumer Products Business,
Tech Mahindra

About the Author

Srinivas Burra is Vice President and Head of the North America Consumer Products business at Tech Mahindra, where he partners with leading CPG companies on growth, technology modernization, and AI-led transformation. His work spans consumer goods, manufacturing, digital commerce, and enterprise platforms, with a consistent focus on turning technology investment into measurable business outcomes. His current focus is how AI is reshaping the software development lifecycle from requirements and code generation through testing, governance, and the operating models needed to sustain them at enterprise scale.

References

1. IBM. (n.d.). AI in the software development lifecycle (SDLC). IBM Think. https://www.ibm.com/think/topics/ai-in-sdlc

2. Deloitte. (n.d.). Agentic AI's impact on software engineering. Deloitte. https://www.deloitte.com/us/en/services/consulting/articles/agentic-ai-impact-on-software-engineering.html

3. ScienceDirect. (2026). Research article on AI and software development. ScienceDirect. https://www.sciencedirect.com/science/article/pii/S1110866526000320?

4. Dev Interrupted. (n.d.). The playbook for your Agentic SDLC. Dev Interrupted. https://devinterrupted.substack.com/p/the-playbook-for-your-agentic-sdlc

5. Tech Mahindra. (n.d.). Telecom giant slashes development time by 83%. Tech Mahindra. https://www.techmahindra.com/insights/case-studies/telecom-giant-slashes-development-time-83-percentage/

About Tech Mahindra

Tech Mahindra (NSE: TECHM, BSE: 532755) offers AI-powered technology consulting and digital solutions to global enterprises across industries, enabling transformative scale at unparalleled speed. With 150,000+ professionals across 90 countries, Tech Mahindra provides a full spectrum of services including consulting, information technology, enterprise applications, business process services, engineering services, network services, customer experience & design, AI & analytics, and cloud & infrastructure services. Tech Mahindra is part of the Mahindra Group, founded in 1945, one of the largest and most admired multinational federation of companies. For more information on how TechM can partner with you to meet your Scale at Speed™ imperatives, please visit https://www.techmahindra.com

www.techmahindra.com
www.linkedin.com/company/tech-mahindra
www.twitter.com/tech_mahindra

Copyright © Tech Mahindra Ltd 2026. All Rights Reserved.

Disclaimer: Brand names, logos, taglines, service marks, tradenames and trademarks used herein remain the property of their respective owners. Any unauthorized use or distribution of this content is strictly prohibited. The information in this document is provided on "as is" basis and Tech Mahindra Ltd. makes no representations or warranties, express or implied, as to the accuracy, completeness or reliability of the information provided in this document. This document is for general informational purposes only and is not intended to be a substitute for detailed research or professional advice and does not constitute an offer, solicitation, or recommendation to buy or sell any product, service or solution. Tech Mahindra Ltd. shall not be responsible for any loss whatsoever sustained by any person or entity by reason of access to, use of or reliance on, this material. Information in this document is subject to change without notice.