

SoftServe



CASE STUDY

How Fortified Health Security Rebuilt Its Entire SDLC Around AI





Executive Summary

Fortified Health Security is a top-ranked MSSP specializing in healthcare cybersecurity. They needed to accelerate software delivery without compromising the strict security, compliance, and quality guardrails that the domain demands. Their traditional SDLC relied on manual execution across requirements, coding, testing, and documentation, which made it difficult to keep pace with evolving regulations and a market moving faster than their delivery model could match.


Their flagship initiative, Central Command, set out to consolidate cybersecurity operations, enhance threat visibility, and improve client engagement. It required a development approach capable of integrating AI-driven capabilities while maintaining strict healthcare security and governance standards.



How did Fortified approach the change?

Fortified implemented an Agentic SDLC framework built on two pillars:

- 1 Context Engineering**
- 2 Spec-Driven Development**



Fortified Health Security went further than approving the initiative, they championed it. Craig Badcock (VP Fortified Health Security) with SoftServe leadership drove both the organizational and workflow changes at every touchpoint, and the shared willingness to embrace and challenge the unknowns on a daily basis defined a strong partnership.

Leadership required the entire 20+ person team to adopt the agentic approach as the default way of working, not an optional experiment. Every new team member is onboarded directly into the framework, making agentic delivery the baseline rather than the exception. This level of client sponsorship has been the single most important accelerator.

The full delivery team (backend, frontend, QC, business analysis, and UI/UX) embraced an experimental mindset. Engineers actively test different AI agents, prompt strategies, and workflow configurations, sharing what works and discarding what doesn't. The team built custom tools, including a RAG-for-Code solution (a retrieval technique that pulls the most relevant parts of the codebase into each AI prompt), to improve the relevance and accuracy of AI-generated outputs within the codebase. Structured training and onboarding programs turned engineers into effective AI operators, while shared context (project rules, templates, and standards embedded directly into AI workflows) ensures consistency across all streams.



What does the agentic framework actually look like?

AT THE CORE OF THE FRAMEWORK:

Agent memory and instructions files are living, version-controlled documents containing architecture overviews, development standards, and AI agent guidelines. They act as a shared context layer for both engineers and AI agents, backed by a Research-Plan-Execute workflow that embeds governance checkpoints before any code is generated. What makes this approach distinctive is the architecture behind it: a single coordination repository acts as the authoritative source for all agent instructions and project knowledge. It aggregates, across multiple repositories, architectural maps, inter-repo dependency graphs, API contracts, business rules, compliance requirements, and domain glossaries into one unified knowledge base.

—

This gives AI agents the same deep contextual awareness a senior engineer builds over months on a project.

EQUALLY CRITICAL IS HOW THAT CONTEXT REACHES THE AGENT:

Rather than flooding it upfront, the framework applies progressive disclosure, staging and revealing context in layers as the agent advances through each phase of the task. The right architectural context surfaces during research, the relevant interface contracts appear during planning, and implementation standards govern execution, ensuring the agent always operates with precision rather than noise.

THIS REDEFINED TEAM ROLES ACROSS THE BOARD:

- Business Analysts moved from manually authoring requirements and specs to directing AI agents to generate user stories and documentation;
- Designers moved from hand-crafting prototypes to orchestrating AI-powered tools that produce UI concepts and production-ready assets;
- Engineers moved from writing every line of code to authoring system-level instructions;
- and QA Engineers transitioned from manual test creation to orchestrating AI-generated test suites.





What did Fortified's agentic SDLC deliver?

Within two quarters (Q4 2025 – Q1 2026) the framework delivered quantifiable outcomes:

~2 FTE

of additional capacity
per quarter

~2 hrs

saved per task on
average

~60%

of the sprint work executed
via agentic approach

100%

agentic SDLC adoption
across all streams

A custom measurement framework tracks AI tool utilization, efficiency gains, and costs, including the ROI multiplier.

Those productivity gains had a compounding effect. The team took on broader scope and onboarded new engineers faster.

“

True leadership was the catalyst for a complete reimagination of the entire SDLC, utilizing AI at every touch point and redefining engineering roles. This was not driven by a single tool, metric, or hype cycle, but a sustained team commitment to redefining how work gets done. SoftServe brought the leadership and engineering rigor for the hard work to make this transformation real. As our agentic SDLC model evolves, it continues to reveal new possibilities, and Fortified Health Security is energized by the pace of progress.

— Craig Badcock,
VP, Fortified Health Security



Where is Fortified headed next?

The next phase of adoption is squarely focused on quality and test automation. The team is training QC and ATQC engineers to leverage the agentic framework to optimize their workflows, structuring specs, test plans, and validation cycles so AI agents can handle the repetitive validation load. At the same time, humans concentrate on edge cases and judgment calls. This will significantly simplify and accelerate the validation stage. Beyond the SoftServe team, the initiative is expanding into Fortified Health Security's own engineering organization, working to embed agentic processes to create a unified, AI-augmented delivery culture.

Asking a 20-person team to abandon their existing workflow and work in a fundamentally different way, immediately and as the default, was a serious ask. Fortified made it anyway. That decision is what separated a proof of concept from a genuine change in how work gets done. This was not driven by a single tool, metric, or hype cycle, but a sustained team commitment to redefining delivery from the ground up.

The next horizon is already in view:



QA automation embedded into the agentic framework



The model expanding into Fortified Health Security's own engineering organization to create a unified, AI-augmented delivery culture across the entire business

How could an agentic SDLC benefit your business?



If your team is still manually executing across requirements, coding, testing, and documentation, it could benefit from an agentic SDLC similarly.

Contact SoftServe's agentic delivery team to explore where an agentic framework can be implemented.


[CONTACT US →](#)



About SoftServe

SoftServe is a premier IT consulting and digital services provider. We expand the horizon of modern technologies to solve today's complex business challenges and achieve meaningful outcomes for our clients.

Social Links



info@softserveinc.com
www.softserveinc.com

Contact

NORTH AMERICAN HQ

201 W 5th Street, Suite 1550
Austin, TX 78701
+1 866 687 3588 (USA)
+1 647 948 7638 (Canada)

EUROPEAN

30 Cannon Street
London EC4 6XH
United Kingdom
+44 333 006 4341