



LM Data & AI Enablement

The Vibe Coding Epidemic

Spec-Driven Development as a Remedy

September 22, 2026

Melissa Bradshaw
Chief Engineer for Data & AI Enablement
Lockheed Martin Enterprise Business & Digital Transformation

Even experienced engineers vibe code... without governance

16

Months

65

Hackathons

2,300

Engineers

4

Business Areas

5

Functions

1

Pattern



AI can introduce insecure patterns, unsafe defaults, missing validation, or unnecessary dependencies that appear functional before they are adequately reviewed: risks that only surface under real data, adversarial input, or production load.



Generated code enters the product faster than teams can understand, test, and maintain it: the bottleneck is no longer typing the code, it is the human capacity to review, verify, and trust what the AI produced before it ships.



The result is AI slop: software that runs in a demo but cannot be confidently verified, operated, or evolved in production. The team owns code they cannot safely change, debug under load, or hand off to the next engineer.



Our product

SEMP^{LM}™ (Systems Engineering Management P^{LM}atform) is an agentic AI platform that performs ISO 15288-compliant systems engineering processes through specialized, orchestrated agents under engineer direction.

Mission: AI performs bounded work while engineers retain accountability for consequential decisions.



ISO 15288 alignment

Work maps to recognizable technical processes.



Process boundaries

An agent cannot operate beyond its assigned engineering boundary.



Inspectable evidence

Outputs preserve source, rationale, and verification path.




Engineer authority

People approve, revise, reject, or escalate consequential work.



SEMP^{LM}™ generates an artifact chain following ISO 15288.

The value is a linked engineering record across technical work—not a stack of generated documents.



Our methodology

Spec-Driven Development (SDD) is an AI skill that embeds SE discipline into AI-assisted development through a five-layer specification architecture with adversarial analysis and security guardrails.

Mission: accelerate implementation while teams retain control of requirements, architecture, and code quality.



Specification alignment

Work maps from approved requirements and design decisions into implementation tasks.



Implementation boundaries

AI cannot modify the product beyond the assigned specification and repository scope.



Reviewable change evidence

Changes preserve the specification, rationale, code diff, tests, and review path.



Engineer authority

People approve, revise, reject, or escalate consequential implementation decisions.



SDD.

SPEC-DRIVEN DEVELOPMENT WORKFLOW

FROM SEMPL · 8 ARTIFACTS

CRM CONCEPT

The specification chain makes AI-assisted implementation reviewable.

SDD has five layers that work together to harness AI generated code.

Layer	Purpose	Owner
Guardrail Specs	Engineering standards per tier (infrastructure, application, agents)	Technical leads
Requirement Specs	Traceable requirements with acceptance criteria	Product owners
Design Specs	Detailed design decisions that satisfy requirements	Developers
Task Lists	Ordered work plans with dependency tracking	Developers
Pipeline Deliverables	Architecture docs, RVTM, SBOM, changelogs	Automated



Adversarial analysis developed with information security organization

A mandatory adversarial analysis phase cross-references all specifications before implementation begins. The analysis systematically identifies:

- Constraint gaps (guardrail requirements with no corresponding specification)
- Security attack surfaces in proposed architectures
- Failure modes with no graceful degradation
- Requirement conflicts and implicit assumptions
- Compliance gaps against mandated standards

Findings are severity-rated (Critical, High, Medium, Low). All Critical and High findings must be remediated before code is written.

Critical

High

Medium

Low



The SDD Alignment Workshop turned a method from imposed process into team-owned practice.

Progressive disclosure and working specifications made the method immediately useful to development work.

Shared vocabulary

Layer concepts around the team's actual work and the next problem ahead.



Working outputs

Produce real product specifications that become inputs to the next development sprint.




Team ownership

The method becomes a practice the team can apply and not just a process imposed on them.




SEMPTM product is demonstrating the methodology's effectiveness.

Multiple benefits emerged from using the spec-first approach.



1. Parallel development

Development happens across agent teams with no integration conflicts; specifications serve as contracts between teams.



2. Complete traceability

The traceability starts with ISO 15288 clauses from requirements to implemented agent behaviors.



3. Rapid results with rigor

Speed comes from clarity, not from cutting corners. Quickly developing features enables actionable feedback.



How is your organization governing AI-generated code?

Systems engineering must evolve for a world in which code generation is nearly free.

Thank you!

melissa.bradshaw@lmco.com