



LM Data & AI Enablement

Speed and Innovation Through Spec-Driven Development

Building an agentic ISO 15288 platform at Lockheed Martin

September 22, 2026

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

Introduction

The context for our case study

Lockheed Martin – Overview

Lockheed Martin is a global defense technology company driving innovation and advancing scientific discovery. Our all-domain mission solutions and 21st Century Security® vision deliver transformative technologies that deter potential adversaries and ensure America and its allies can achieve peace through strength.

123K

Employees
worldwide

93%

Employees located
in the U.S.

20%

Employees who
are veterans

72K

Engineers, scientists
and technologists

340+

Facilities
worldwide

100+

Years of
innovation



Lockheed Martin – Business Areas

Our corporation, headquartered in Bethesda, Maryland, is organized around our core business areas specializing in defense tech.



Aeronautics, with approximately \$30.3 billion in 2025 sales, which includes tactical aircraft, airlift, sustainment and aeronautical research and development lines of business.



Missiles and Fire Control, with approximately \$14.5 billion in 2025 sales, which includes the High Mobility Artillery Rocket System (HIMARS) and PAC-3 Missiles as some of its high-profile programs.



Rotary and Mission Systems, with approximately \$17.3 billion in 2025 sales, which includes Sikorsky helicopters, maritime systems, sensors, radar systems, command and control, combat simulation and training, undersea systems and full-spectrum cyber capabilities.



Space, with approximately \$13.0 billion in 2025 sales, which includes development of national security, civil, commercial, commercial and weather spacecraft, strategic and missile defense solutions and hypersonics missions.

Without trustworthy AI engineering, speed creates both engineering risk and maintenance debt



Plausible but incorrect AI output can become an engineering baseline, allowing a bad need, requirement, or design assumption to propagate downstream before it is discovered.



When teams cannot show where an AI-assisted decision came from, reviewers cannot determine whether the result is justified, governed, or defensible.



Generated artifacts can look complete while concealing missing constraints, untested interfaces, conflicts, and unaddressed stakeholder needs.



AI-generated code can work locally while violating the product's architecture, module boundaries, established conventions, and implicit interface contracts.



AI can introduce insecure patterns, unsafe defaults, missing validation, or unnecessary dependencies that appear functional before they are adequately reviewed.



When generated code enters the product faster than teams can understand, test, and maintain it, the result is AI slop: software that is difficult to debug, modify, verify, and safely operate.

Agenda

- **SEMP^{LM}™ (Systems Engineering Management Platform)**

An agentic AI platform that performs ISO 15288-compliant systems engineering processes through specialized, orchestrated agents under engineer direction.

- **Spec-Driven Development (SDD)**

An AI skill and methodology that embeds SE discipline into AI-assisted development through a five-layer specification architecture with adversarial analysis and security guardrails.

- **A workforce transformation model**

The deliberate, tiered investment required to move teams and organizations from traditional workflows to AI-enabled SE practice.



SEMPLE™ agents perform ISO 15288 processes

Agentic AI executes SE processes under engineer orchestration with
built-in trust

Solution Origin

SEMPLE™ began with a systems-engineering question:

Can AI perform 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} begins with a mission need, not a chat prompt.

CRM concept illustration: a customer relationship management product enters through an explicit mission, vision, and constraints.

Mission

Give sales teams a unified, governed customer view.

Vision

From customer need to an engineered, traceable CRM capability.

Constraints

Auditability • least privilege • validated authority.




Each agent has an explicit contract: process boundary, artifact, and quality assessment.

SEMPLE™ is intentionally more constrained than “ask an agent to solve the problem.”



1. Process boundary

The agent is assigned to one technical process; it does not silently expand authority.



2. Expected artifact

The work produces a named engineering artifact—not an unstructured conversation.




3. Quality assessment

Defined criteria assess the artifact and return an explicit disposition.



SEEMPL™ persists an artifact chain through the verification milestone following ISO 15288.

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



Three execution patterns preserve different kinds of evidence and control.

Choose the smallest execution structure that preserves the process need.

Linear agent

One Strands agent + structured output.
Use for bounded evaluation or review.



Graph agent

Typed nodes share invocation state.
Use when intermediate artifacts and state must be explicit.




Stage orchestrator

Tool-driven coordination across agents.
Use when feedback, iteration, and escalation control the next step.




Trust is engineered as a system of controls.

The product limits what the model is trusted to do, then validates, records, tests, and escalates the rest.



SEMPLE™ reusable trust harness evaluates behavioral properties instead of comparing one response to a golden paragraph.



SDD embeds SE discipline into AI development

The spec-first methodology governs how we build AI systems through adversarial analysis and security guardrails

Solution Origin

SDD began with a software-delivery question:

Can AI 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.





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




SEMPTM product is developed using SDD, 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.



Building the human capital

A technical architecture becomes useful only when teams can apply it with judgment.

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.



Traditional Hackathons

- ~65 events, ~2300+ participants
- Immersive multi-day events where teams or organizations learn AI skills, upskill practitioners, and explore use cases relevant to their domains.
- Events span business areas including ERP, Manufacturing, F-35, HR, Cyber, Space, and international (UK, Canada).



Hackathons identify high-apptitude participants

Build Events

- 4 completed
- Focused engagements targeting a single use case, culminating in a production-deployed solution.
- Unlike hackathons, these are scoped to deliver operational capability, not just exploration.



Build events prove production viability



Frontier Operations

- 3 cohorts of ~50 builders
- An intensive one-week bootcamp for high-aptitude technical talent focused on pure tools and methodology immersion at the human-AI boundary.
- Skills include Boundary Sensing (where AI capability ends and human expertise begins), Seam Design (architecting human-AI transition points), and Failure Mode Awareness.
- Graduates serve as a feeder pipeline for teams supporting enterprise AI demand.



FrontierOps creates advanced practitioners



Domain experts are part of the trust architecture.

Partnership with security, quality, and mission SMEs validates whether output is correct, compliant, and useful.

Security partners

Challenge specifications and architectures through adversarial analysis.

Engineering SMEs

Validate real process, domain constraints, and evidence needs.

Product / mission owners

Judge whether work solves the right problem and is ready for the next decision.



Closing message

The transferable contribution is a product + method + workforce model.

Trustworthy Innovation

SEMPL™ Governs

Specialized agents perform bounded ISO 15288 work through reusable orchestration patterns.

SDD Executes

Specifications and adversarial analysis make AI-assisted implementation reviewable and traceable.

Workforce Enables

Workshops, build events, and advanced practitioner development turn a method into practice.



Thank you!

melissa.bradshaw@lmco.com