Expert-Driven Development

Open-Source AI-Assisted Development Training for Department of War Personnel

SSgt Jesse C. Morgan, USMC

Marine Corps Detachment, Presidio of Monterey

Defense Language Institute Foreign Language Center

January 2026

The Problem: Expertise Without Tools

Domain experts understand their problems better than anyone.

But they can't build the solutions themselves.

- Traditional development: contractors, months of lead time, six-figure budgets
- IT backlogs: requests sit for months or years
- By the time a tool ships, the problem has changed
- The people closest to the problem remain disconnected from the solution

Meanwhile: 80% of workers given AI tools quit within weeks (Microsoft). Access without training is worse than no access.

80%
abandon AI tools within weeks

-19 pp
accuracy when untrained users hit the jagged frontier

$0
EDD tool development cost

The Solution: EDD in 60 Seconds

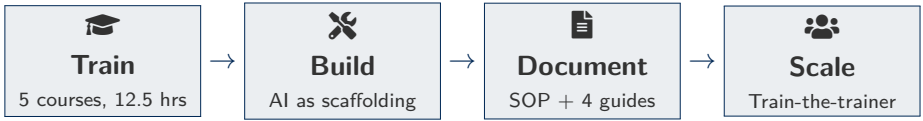

What EDD is:

- Training methodology for domain experts
- Governance framework (SOP) for responsible building
- AI is scaffolding — removed after construction
- Tools are standard apps on standard infrastructure
- Open-source, MIT-licensed, UNCLASSIFIED

What EDD is not:

- Not a coding bootcamp
- Not an AI product or platform
- Not dependent on any single AI tool
- Not a replacement for IT professionals
- Not a shortcut around governance

Built on Research, Not Intuition

Every design decision traces to peer-reviewed empirical findings — 7 studies, 300K+ subjects.

Study	Finding	EDD Design Impact
Microsoft Work Trend Index (300K+)	80% abandon AI within weeks	Management framing, not tool training
BCG-Harvard 2023 (758 consultants)	+40% quality inside frontier; -19pp outside	Frontier Recognition; Centaur/Cyborg modes
Stanford/MIT 2025 (5,172 agents)	Novices +34%; experts flat	Quality Judgment as core skill
Mollick 2026 (Wharton MBA)	Management skills predict AI effectiveness	"Learning to manage, not use"
UK Gov 2025 (20K employees)	25 min/day saved; 80%+ retained	Gov orgs benefit with proper training
Stanford DEL 2025	Entry-level postings down 35%	Apprentice Problem protocol
OpenAI GDPval (1,320 tasks)	AI at expert parity on 48% of tasks	Focus on judgment over prompting

The Six 201-Level Skills

These are **management skills**, not prompting techniques. Every module maps to at least one.

1. Context Assembly

Curate background, constraints, and examples *before* tasking AI. Prevents over-sharing sensitive data.

2. Quality Judgment

Spot reliable vs. unreliable content within the *same* output. The Red Pen Review exercise.

3. Task Decomposition

Break work into AI-appropriate chunks. Decide what to delegate vs. retain.

4. Iterative Refinement

First output is a starting point, not a deliverable. Structured revision passes.

5. Workflow Integration

Embed AI into recurring processes with documented playbooks. Human/AI labels.

6. Frontier Recognition

Know where AI excels and fails for *your* work. Share failure cases across the org.

Training Curriculum: 6 Courses, Every Role

	Course	Hrs	Audience	Outcome
1	AI Fluency Fundamentals	2	All personnel	Six 201 skills, jagged frontier, when to trust AI
2	Builder Orientation	2	Aspiring builders	Build a working prototype from a real problem
3	Platform Training	4	Builders	Build 3 tools on Power Platform; centaur/cyborg
4	Advanced Workshop	4	Experienced	Frontier mapping, verification, teach others
5	Supervisor Orientation	0.5	Leadership	Permission culture, evaluate proposals
6	Full-Stack Capstone	8	Advanced	Go + React + AI chat; personal time

Universal Pathway: Course 1 → Course 5 (leadership)

Builder Pathway: Course 1 → 2 → 3 → 4 → 6

Instructor Certification: Complete as student → shadow → co-teach → certified. Each unit maintains 2+ instructors. PCS-resilient.

Course 1–2: From Zero to Prototype

Course 1: AI Fluency Fundamentals

2 hours | Required for all personnel

- Opens with the 80% abandonment problem
- Introduces all six 201-level skills
- **Red Pen Review**: verify AI output before signing
- Centaur vs. Cyborg work patterns
- Delegation Equation (Mollick)
- Verification hierarchy: high/medium/low stakes
- Closes with personal AI action plan

No tools required. No Power Platform. Pure judgment skills.

Course 2: Builder Orientation

2 hours | Prerequisite: Course 1

- "Come with a problem you want to solve"
- Live instructor demo: real tool in 30 min
- Students build their first prototype
- Task Decomposition in practice
- Iterative Refinement: 70% $\rightarrow$ 95%
- Debugging with AI assistance
- Failure sharing: what went wrong?

Students leave with a working prototype of their own idea.

Courses 3–6: Building to Mastery

Course 3: Platform Training

4 hours | Build 3 tools

- Build #1: Guided (follow along)
- Build #2: Semi-guided (prompts only)
- Build #3: Independent (your problem)
- Power Platform: Power Apps, Power Automate, SharePoint
- Mandatory failure-sharing session
- Begin frontier mapping

Course 4: Advanced Workshop

4 hours | Organizational capability

- Domain-specific frontier maps
- Verification protocols & QA checklists
- Workflow playbooks (Human/AI labels)
- The Apprentice Problem
- 60-Second Teach exercise
- Published deliverables for org

Course 5: Supervisor Orientation

30 min | Highest leverage

- Permission culture vs. fear culture
- Evaluating AI-assisted work
- Protecting junior development
- One supervisor enables an entire section

Course 6: Full-Stack Capstone

8 hours | Personal dev time

- Go + React + Docker + AI chat
- Commercial frontier tools
- Unclassified, off-network only
- Readiness investment

Standard Operating Procedure: 8-Phase Workflow

Phases 1 and 6 are governance gates (supervisor signature / independent QA review)

Governance Gates:

- **Supervisor signature** required before development begins
- **Compliance checklist** (Phase 2): data classification, PIA, platform
- **QA Review** (Phase 6): independent reviewer, 15-item checklist
- CDAO RAI Toolkit alignment (QA Item 8)

Six Defined Roles:

- **Developer** — builds the tool
- **QA Reviewer** — independent gate
- **Supervisor** — approves & oversees
- **Program Coordinator** — registry, training, scaling
- **Cybersecurity Office** — security guidance
- **Privacy Officer** — PIA review

Documentation Package: User Guide, Replication Guide, Adaptation Guide, Maintenance Guide.

A new person can maintain the tool using only the documentation.

Approved Tools & Platforms

AI Assistants:

Tool	Level	Best For
GenAI.mil	IL5 (CUI)	Primary portal
CamoGPT	IL5/IL6	Classified + API
M365 Copilot	Enterprise	Docs, email, Excel
Copilot Studio	Enterprise	Bots, workflows
Azure OpenAI	FedRAMP High	Custom dev
ChatGPT	Unclass only	Via GenAI.mil
Gemini	Unclass only	Via GenAI.mil

Deployment Platforms:

- Power Platform (Power Apps, Power Automate)
- SharePoint Online
- Azure Government
- GitHub (public, open-source)

Security Decision Matrix

CUI data?	GenAI.mil
Classified?	CamoGPT (SIPR)
Power App + AI?	Copilot / CamoGPT
General coding?	GenAI.mil
Correspondence?	M365 Copilot

Never

- CUI in commercial ChatGPT/Gemini
- PII without completed PIA
- Classified outside CamoGPT SIPR

Toolkit & Resources

Interactive Toolkit

- Searchable, filterable prompt database
- Organized by development phase, platform, complexity
- Copy-ready prompts with bracketed placeholders
- Live at the EDD site — no install required

Prompt Library

- Curated prompts by use case:
  - Admin tools (equipment tracker, leave system, room booking)
  - Training tools (quiz generator, attendance tracker)
  - Data analysis & reporting
- Designed for M365 Copilot, GenAI.mil, Azure OpenAI

Templates (SOP Appendices)

- Problem Definition Worksheet (Phase 1)
- Development Journal (Phase 2)
- QA Review Checklist (Phase 6)
- Tool Registry Entry (Phase 7)
- PIA Threshold Analysis (Phase 2)
- Compliance Checklist (Phase 2)

FAQ

- "Do I need to know how to code?" — **No.**
- "How long does a first tool take?" — **50–90 hours**
- "What platform should I use?" — Power Platform
- "What about classified?" — Never in unapproved tools
- "Do I need supervisor approval?" — **Yes. Always.**

Full DoD AI Compliance

EDD has been crosswalked against every major DoD AI governance document. **No gaps.**

Framework	Status
DoD RAI 5 Principles	Covered
CDAO Guidelines & Guardrails (Jul '24)	Covered
Sec 1: Risk Criteria	Covered
Sec 2: Suitability/Feasibility	Covered
Sec 3: Training Data Quality	Covered
Sec 4: Cyber Resilience	Covered
Sec 5: Traceability	Covered
Sec 6: Harms/Impacts/Risks	Covered
Sec 7: Human-Machine Teaming	Covered

Framework	Status
AI Cyber RM Tailoring Guide (Jul '25)	Covered
DON GenAI Guidance	Covered
FY2026 NDAA AI Provisions	Covered
DoW AI Strategy (Jan '26)	Covered
MARADMIN 018/26	Covered

Section 7 is the Knockout

The CDAO memo *recommends* GenAI training on risks, effective use, and capability limitations. EDD *delivers* exactly that. Five courses. 12.5 hours. With exercises, QA protocols, and published deliverables.

Course 6 uses commercial frontier tools on personal time/devices, unclassified only — fully within the commercial-tools-for-unclassified lane.

Scaling Model & Metrics

90-Day Pilot Model:

1. **Days 1–15:** Program Coordinator designated. SOP adopted. Training scheduled.
2. **Days 16–30:** All-hands AI Fluency (Course 1). Identify 3–5 pilot builders.
3. **Days 31–60:** Builder courses (2–3). Pilot developers build first tools.
4. **Days 61–75:** QA reviews, remediation, deployment.
5. **Days 76–90:** Assessment, metrics review, expansion decision.

Expansion:

Pilot devs become QA reviewers & mentors.

Certified instructors PCS with the capability.

Tool Registry grows organically.

Program Metrics:

Metric	Target
Course 1 completion	$>90\%$
Builder completion	$>70\%$
Tools deployed/quarter	≥ 2
Active users per tool	>10
User satisfaction	$>4.0/5$
Doc compliance	100%
Time to first deploy	<90 days

Success Criteria:

- ≥ 1 tool deployed within 90 days
- $\geq 80\%$ Course 1 completion
- ≥ 2 certified instructors
- Positive user feedback

“The guardrails don’t work if nobody knows how to stay inside them. That is what EDD solves.”

jeranaias.github.io/ExpertDrivenDevelopment

github.com/jeranaias/ExpertDrivenDevelopment

jesse.morgan@dliflc.edu

Open-source | MIT Licensed | Ready to deploy today