

# Your people are waiting for you to say yes.

A leadership briefing on permission culture, evaluating AI-assisted work, and the apprentice problem.

WEEK 5

LEADERSHIP BRIEFING

30 MIN · OR EXTENDED JOINT SESSION

# Who this week is for

## WEEKS 1-4 · BUILDERS

### Marines learning to use AI and build tools

- Fluency, prototyping, Power Platform, advanced workshops
- Hands-on by design — students were the workers
- Time was spent *building*

## WEEK 5 · LEADERSHIP

### Supervisors and decision-makers

- Section leaders, OICs, SNCOs, civilians who approve work
- You don't need to build — you need to *evaluate* and *permit*
- Time is spent on culture, judgment, and three decision exercises

# Pick the mode that fits today

## MODE A · 30 MINUTES

### Leadership briefing

The standalone briefing. Five short modules. Tight script. Closes with one specific ask.

- 5 min · Why this matters now
- 5 min · Permission culture
- 10 min · Evaluating AI-assisted output (2 scenarios, light debrief)
- 5 min · The apprentice problem
- 5 min · Quick reference + close

## MODE B · 60-90 MINUTES

### Joint session with builders

Same slides. Speaker notes branch into discussion prompts and a third decision exercise.

- Each module pauses for floor discussion
- Scenarios A and B run as paired-talk + report-out
- Adds Exercise C (Spot the Apprentice Risk)
- Builders share one real artifact for live evaluation

# Five modules · 30 minutes (or 60-90 in joint mode)

Time	Module	Focus
0:00 - 0:05	1. Why this matters now	The permission gap, the data, the directive
0:05 - 0:10	2. Permission culture	What "yes" sounds like; what kills adoption; guard rails (not roadblocks)
0:10 - 0:20	3. Evaluating AI-assisted output	Four questions; <b>Exercises A &amp; B</b>
0:20 - 0:25	4. The apprentice problem	Junior development, three supervision checks; <b>Exercise C</b> (joint)
0:25 - 0:30	5. Quick reference & commit	Reference card, leadership commitment, Week 6 preview

**Your people are **already**  
using AI.  
They're waiting for you to  
say **yes**.**

SOURCE: MOLLICK, *ONE USEFUL THING* · ENTERPRISE FIELD RESEARCH

**MODULE 1 · 5 MINUTES**

# **Why this matters now**

The data, the directive, and what it means for your section.

# Five reasons "later" is no longer an option

**~80%**

workers abandon AI within  
weeks *without* support  
**MICROSOFT**

**25 min**

/day saved *with* proper support  
**UK GOV · 20K EMPLOYEES**

**2026**

"Year of Military AI Dominance"  
**DOW AI STRATEGY · JAN  
2026**

**49B**

new dedicated AI/ML officer  
career field  
**ARMY · OCT 2025**

**018/26**

MARADMIN designating  
GenAI.mil enterprise  
**MARINE CORPS**

The pattern: **everywhere AI deployment was paired with leadership support, it stuck.**  
Everywhere it wasn't, it died.

*Sources: Microsoft Work Trend Index · UK GDS 2025 · DoW AI Strategy Jan 2026 · Army HRC · MARADMIN 018/26.*

**It's not optional anymore.**

Your command directed it. The DoW directed it. The other services are already past the starting line. Your section is either reporting AI activity or reporting why it isn't.

MODULE 2 · 5 MINUTES

# Permission culture

What "yes" sounds like · what kills adoption · guard rails, not roadblocks.

# The language your section is listening for

## WHAT "YES" SOUNDS LIKE

---

- "Try it. Show me what you build."
- "I don't need to understand *how* it works. I need to understand *what* it does."
- "If it saves time and meets quality standards, we should do it."
- Protected time for learning — even 30 minutes a week.
- Public recognition when somebody builds something useful.

## WHAT KILLS ADOPTION

---

- "I need to approve every AI interaction."
- "Don't use AI for anything official."
- "We'll wait until there's a formal policy."
- Punishing experimentation failures.
- Treating AI use as suspicious or lazy.

# What supervisors should *not* do

**1****Don't ban AI use outright.**

Prohibition drives it underground where you can't guide it. You lose visibility and you lose control.

**2****Don't require approval for every AI interaction.**

Creates a bottleneck that kills velocity. Review the output, not every query.

**3****Don't assume AI output is automatically wrong.**

Reflexive skepticism kills adoption. Verify like any work product — evidence-based, not assumption-based.

**4****Don't skip verification because the Marine is experienced.**

Seniority does not eliminate hallucination risk. Everyone verifies.

# Guard rails, not roadblocks

## The four boundaries that always apply

- All tools go through the **EDD SOP** before deployment.
- AI-generated content for official use must be **human-reviewed**.
- **Sensitive or classified information** never enters unauthorized AI systems.
- Failed experiments are **shared as learning** — not punished as waste.

### Default posture

**"Yes, with appropriate review."**

That single sentence is your operating policy until something specific tells you otherwise.

MODULE 3 · 10 MINUTES

# Evaluating AI-assisted output

Four questions · two practice scenarios · the standard you already use.

# Four questions — every time

---

## 1 Does it work?

Can you demonstrate it doing what it claims to do — live, end to end?

## 2 Is the output accurate?

Are facts, references, citations, numbers, and dates verifiable against source material?

## 3 Does it follow the SOP?

Has it gone through the proper review path (peer, SME, security as required)?

## 4 Does it save time?

Is there a credible before/after comparison — and is the QA time reasonable for the gain?

You do *not* need to evaluate *how* AI produced it. The same standards as any work product apply: **if a Marine handed this to you without mentioning AI, would you sign it?**

# AI-generated leave policy summary

## SCENARIO

A Marine presents you with a one-page summary of leave policy for Emergency Leave Authorization. The summary is clear, well-formatted, and cites **MCO P1050.3K**. The Marine tells you it was generated using ChatGPT and reviewed for accuracy. They want to distribute it to the section.

## YOUR MOVE

What do you ask first? What gets the green light? What earns a "not yet"?

# What a strong supervisor asks

- **Can you verify the MCO citation?** Pull up the actual order. Does the summary match?
- **What did you give the AI as input?** Full MCO text, or general questions? Context quality drives output quality.
- **Who reviewed this for accuracy?** AI-generated policy needs SME verification — not just proofreading.
- **What's the update path?** If policy changes, who refreshes this — or do outdated summaries circulate forever?

## The principle

The Marine did the right thing — they brought it to you. Your job is to verify the *output* and the *process*, not interrogate the tool. Approve, ask for SME sign-off, or send it back with a specific gap.

# Automated reporting tool

## SCENARIO

A Marine built a Python script that pulls data from a shared spreadsheet and auto-generates the weekly operations summary. **It runs in 30 seconds. The manual process took 90 minutes.** The Marine wants approval to use it for official reporting.

## YOUR MOVE

The time savings are real. The risk is also real. Which questions go first?

# What a strong supervisor asks

- **Can you demonstrate it with live data?** Show input, process, output. Walk through one full cycle in front of me.
- **What happens when input data is malformed?** Does it fail gracefully — or does it generate confident, incorrect reports?
- **Who else has reviewed this code?** Per the EDD SOP, automation tools need peer review and security review before operational use.
- **How do we verify accuracy each week?** Sampling checks, output validation — what's the ongoing QA?

## The trap

"It saves 89 minutes" is not a Yes. It's a reason to take the four questions *more* seriously, not fewer. The most dangerous AI-assisted work is the kind that's obviously valuable.

MODULE 4 · 5 MINUTES

# The apprentice problem

The risk you can't see this year — and that no one else will catch.

"That training pipeline that was always implicit has broken — and it has to be reconstructed."

— ETHAN MOLLICK · WHARTON, 2025

**–35%**

Entry-level job postings in AI-exposed occupations (*Jan 2023 → Jun 2025*)

**–13%**

Employment among workers age 22–25 in those fields

**5–10 yrs**

until organizations that cut junior roles face an expertise shortage they cannot hire out of

# Three supervision checks for junior development

## 1 Can the Marine explain the output without referencing AI?

If they can't walk you through the logic, they didn't learn — they copied.

## 2 Require periodic tasks completed without AI.

Preserves baseline competency. If the tool disappears, can they still operate?

## 3 Use AI output as a teaching tool.

"Is this correct? How would you verify it? What would happen if this number were wrong?" That's where judgment is built.

**Goal:** AI-augmented Marines, *not* AI-dependent ones.

# Spot the apprentice risk

## SCENARIO · JOINT DISCUSSION

A new SgtSec NCO is assigned to draft initial counseling statements for three junior Marines. They use ChatGPT to produce all three drafts in under five minutes. The drafts are technically correct. They sign and submit them. Their supervisor approves them without comment.

### What was lost

- The new NCO never wrestled with how to characterize a Marine's performance.
- They didn't develop the language a counseling statement actually requires.
- They learned that AI handles people work — and that the supervisor doesn't notice.

### What a permission-culture supervisor does

- Approves AI for the *format and structure*, not the judgment.
- Requires the NCO to write the assessment paragraphs from scratch.
- Reviews together: "What did the AI miss about this Marine?"

MODULE 5 · 5 MINUTES

# Quick reference & commit

A card for your screen · the tools · further reading · Week 6 preview · the ask.

# Supervisor quick reference: AI-assisted work

## 3 QUESTIONS BEFORE YOU APPROVE

- **Demonstrate it.** Live demo with real inputs.
- **Verify accuracy.** Specific checks, not general trust.
- **Who reviewed it?** Peer · SME · security per the SOP.

## 3 SIGNS OF HEALTHY AI USE

- The Marine can explain the output *without* referencing AI.
- Failed experiments are shared openly — learning culture.
- AI is saving time on routine tasks, not replacing judgment tasks.

## 3 WARNING SIGNS

- Marine cannot answer basic questions about their own output.
- AI use is hidden or apologized for — sign of a fear culture.
- Junior Marines are skipping foundational tasks entirely.

**Default posture: "Yes, with appropriate review."**

# Tools your Marines will use

## GenAI.mil

**Default for DoD networks.** IL5-authorized for CUI per MARADMIN 018/26. Hosts Gemini, Grok, ChatGPT inside the DoD boundary. ~1.1M users.

## CamoGPT (Army)

Supplementary capabilities: API access, tool calling, file upload, shared workspaces, IL6/SIPR access for classified network use.

## Commercial: ChatGPT, Gemini

Approved for **unclassified** work only via vendor websites. Supplement DoD tools when CUI is not involved. Also available inside GenAI.mil.

## Always

PII must be anonymized on every platform unless a PIA explicitly authorizes it. Sensitive or classified data never enters unauthorized systems. See the *EDD Approved Tools* page.

# Two briefs to hand your CO

## EDD Executive Brief

15-slide Teams briefing. The "what is EDD, why now, what we're asking from leadership" deck — designed for an O-5/O-6 audience.

[docs/pdf/EDD\\_Executive\\_Brief.pdf](docs/pdf/EDD_Executive_Brief.pdf)

## EDD RAI Compliance Brief

One-page compliance map: how the EDD program aligns with DoD AI guidelines and Responsible AI guardrails. The document for your staff judge advocate or compliance officer.

[docs/pdf/EDD\\_RAI\\_Compliance\\_Brief.pdf](docs/pdf/EDD_RAI_Compliance_Brief.pdf)

If you only forward one thing after this session: **send the Executive Brief to your O-5/O-6 with one sentence — "We are doing this; here is the structure."**

## The bonus capstone

- **Audience:** advanced builders only (elective — has prerequisites)
- **Duration:** 8 hours, builder-track
- **Outcome:** direct AI to write a complete Go + React + Docker application
- **Pattern:** full Cyborg-mode at full-stack scale
- **Prereq:** Advanced Workshop completed + at least one deployed tool

### What this means for supervisors

Most of your Marines will *not* take Course 6 — and that's by design. It's for the small group whose problems exceed the Power Platform envelope.

If a builder in your section asks about it, the question to ask is: *"Show me the deployed tool you built first."*

## LEADERSHIP COMMITMENT

**The structure exists. Your people have the motivation. Your job is to give the permission.**

### WITHIN THE NEXT WEEK

Ask one person in your section what they'd build if they had permission. Listen to the answer. Then say *yes, with appropriate review*.