accenture

DAFITC 2026 / TRACK 4 / WEDNESDAY AUGUST 26

# Accelerating Mission Readiness with AI-Enabled Training

Building AI-fluent Airmen and Guardians at the intersection of intuition and machine intelligence

**Michael Forrester**

AI for Leadership and Organizations Business Stream Lead, Accenture LearnVantage

accenture Federal Services

WHERE YOU ALREADY ARE

accenture

# You have already started. The question is what comes next.

**1.7M**

nearly 1.7 million users on GenAI.mil S18

**100,000+**

custom agents already built by your own people S18

**Agent Designer**

shipped March 2026, for after-action reports and multi-step work S18

THE REQUIREMENTS

accenture

# What you are being asked to do, as of this summer

### NSPM-11, June 2026

Four pillars for the national security enterprise. Directs a DoDD 3000.09 update, standardized test and evaluation, an AI curriculum. S19

### DoW AI strategy, January 2026

AI-first warfighting force, seven Pace-Setting Projects, aggressive data access across classification levels. S20

### DAF AI Strategy and AI Talent Plan, April 2026

Baseline AI literacy across the entire force. Competency becomes everyone's requirement. S21 S22

### SWFT and continuous authorization

Continuous machine-readable evidence replaces the paper package. Assurance, at speed. S8 S9

> **Four directives. One workforce question at the core of all of them.**

THE MANDATE

accenture

# The Department of the Air Force has already decided AI fluency is a mission requirement, not an experiment

THE PLAN

**DAF AI Talent Plan**

April 2026 S1 S22

THE TARGET

**AI Competitive by 2027**

Set as a hard target S2

THE WORKFORCE

**AI specialists (depth)**

**AI-enabled operators (breadth)**

A dual-track model

THE METRIC

**Demonstrated competency**

Not course completion or self-assessment

THE MEASUREMENT GAP

accenture

# Competence with AI is the one thing people cannot self-assess

MEASURED PERFORMANCE

**In a randomized controlled trial, experienced practitioners using AI tools were measurably slower** S17

SELF-REPORT

**They believed they had been substantially faster**

**Self-report and performance moved in opposite directions**

> **This is why the DAF plan measures demonstrated competency, not course completion** S1

DEMONSTRATED COMPETENCY TO OUTCOMES

accenture

# Faster time-to-competency and higher operational confidence

## COMPETENCY DEMONSTRATION

### Training

Knowledge

**Proven by completion**

### Skill

Applied knowledge, in context, with tools

**Proven by assessment**

### Outcome

Applied to a real workflow, producing measured value

**Proven by a working capstone project**

## THE MODEL S38

**AI-MAP, the AI Maturity Appraisal and Roadmap Model, developed with Carnegie Mellon University.** It measures how an organization adopts and delivers AI with predictable, measurable outcomes, and returns a maturity marker plus a roadmap.

## THE MEASUREMENT METHOD

### Assess

Baseline the operator's demonstrated capability before training

### Roadmap

Role path plus coaching at the point of need, grounded in doctrine

### Reassess

Re-measure against the same standard to show the delta

DEMONSTRATION ONE | OUTTHINK

# Where am I on the ladder, and what do I do next?

The question: what do I need to know to be skilled at AI.

accenture

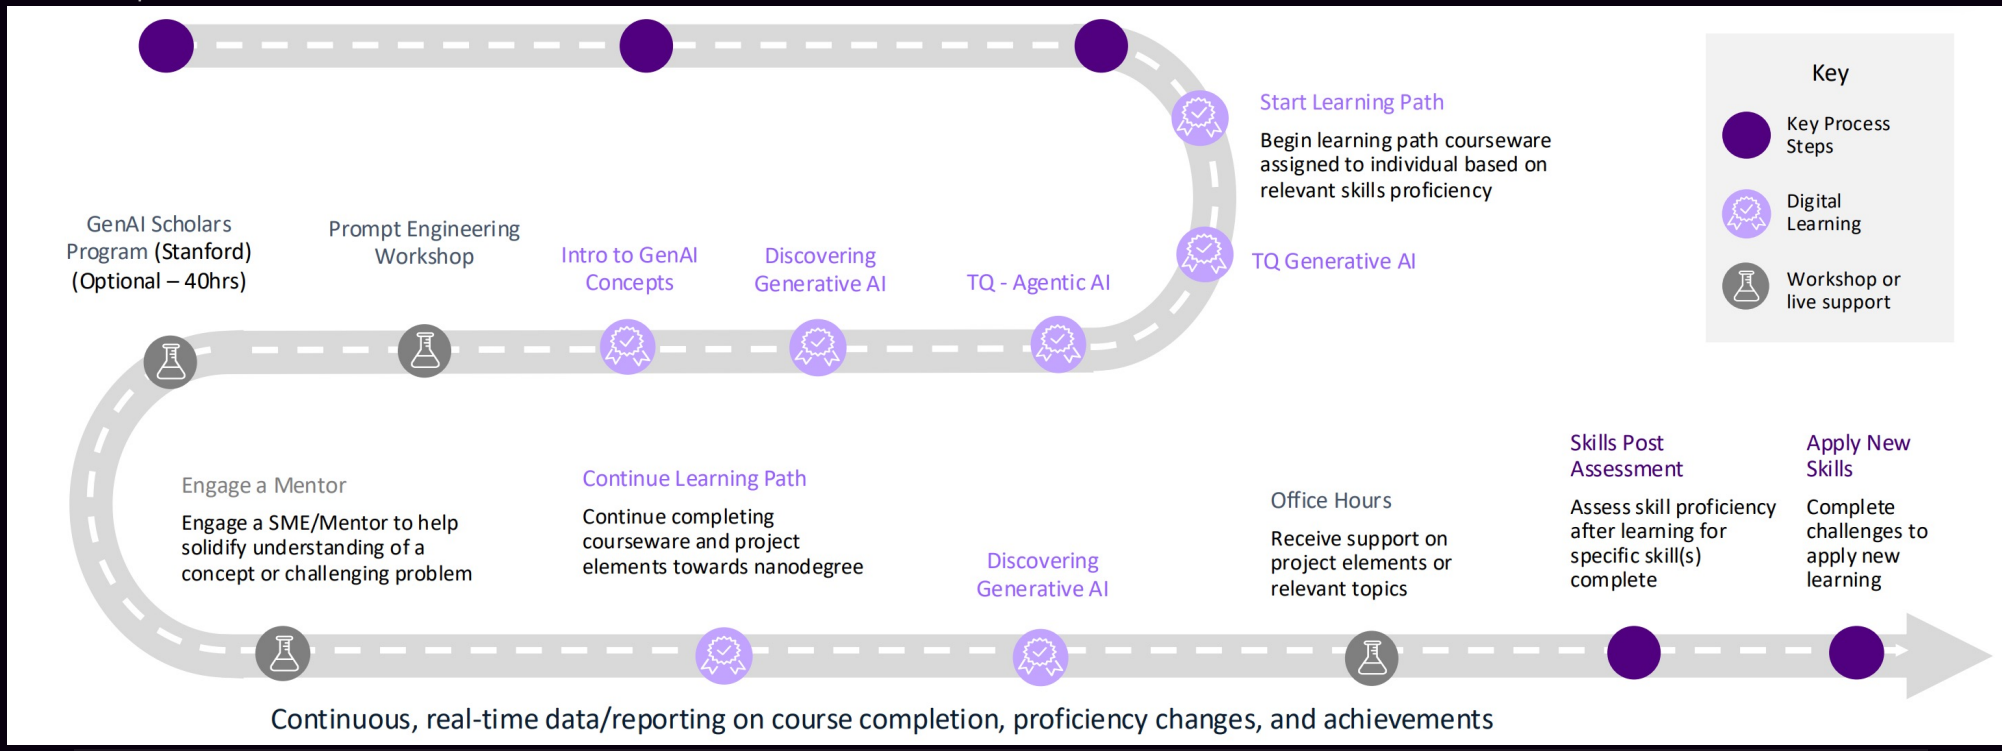


01 Assess against the rungs

02 A personalized path, doctrine-grounded coaching at the point of need

03 Reassess against the same instrument

THE LADDER

accenture

# Five rungs from assistant to action, and most of the value is not at the top

Rung 01

**Assistant**

Ask and answer. No access to your systems.

Rung 02

**Retrieval-grounded**

Answers from your cleared corpus, with citations.

Rung 03

**Tool-using**

Calls systems under supervision, one action at a time.

Rung 04

**Read-only autonomous**

Runs continuously, observes and briefs, changes nothing.

Rung 05

**Action-taking**

Alters systems of record.

> **Rung 4 is the high-autonomy, low-agency coordinate: continuous work, no committed writes to authoritative records. Highest value per unit of risk.** S28

Rung 5: highest barrier. Requires everything below it to be proven first.

THE COMPETENCY

accenture

# What you have to be good at changes as you climb

THE SKILL S35

- **Delegation.** Knowing which work to hand over and which to keep.
- **Description.** Specifying the task precisely enough to be answerable.
- **Discernment.** Judging whether the output is actually right.
- **Diligence.** Verifying, and owning the result.

HOW IT SCALES

Discernment and diligence are the two that have to grow fastest, because they are the only defense against the measurement gap.

**At rung one a bad answer costs you a minute. At rung four it costs you a decision.**

> Experienced operators shift from approving every step to monitoring and intervening. In practice roughly three quarters of agent tool calls stay human-in-the-loop, and less than one percent are irreversible. S29

WHAT TRAINING REQUIRES

accenture

# Training only transfers when it is doctrine-aligned and available at the point of need

Generic AI training does not survive contact with the mission

## 01

### Doctrinal grounding

Guidance must be grounded in your doctrine and tactics, techniques, and procedures (TTP)

## 02

### Point-of-need delivery

Delivered in the flow of the task, not in a classroom weeks earlier

WHY ADOPTION FAILS

accenture

# Pilots do not fail because the models are weak

01 The failure mode is workflows that were never redesigned S23

02 Buying or partnering succeeded roughly twice as often as building internally S23

03 Back-office work returned more than the sales and marketing pilots that absorbed most of the budget S23

04 Gartner forecasts more than 40 percent of agentic projects canceled by end of 2027 S24

**WORKED**

PANDA, Air Force predictive aircraft maintenance: a bounded workflow, verifiable output, a maintainer in the loop. S37

**DID NOT WORK**

Programs that could not use their own model outputs because data rights were never negotiated, with no lessons written down. S25

THE FIRST WORKFLOW

accenture

# The first one is a selection problem, not a technology problem

01

**A bounded system, not open-ended**

02

**A verifiable output, tested against a measured baseline**

03

**Low cost of a wrong answer**

04

**Cheaper or faster than the human doing it today**

05

**A human positioned to course-correct**

06

**Data access, classification, and compliance settled before the build**

07 **Are the fundamentals handled: is this workflow already automated? If yes, strong AI candidate. If no, AI plus a human partner might get you to baseline.**

THE FORMAL VERSION These gates are the field version of a structured use case prioritization framework: rank candidate use cases on value, feasibility, and organizational readiness, then sequence deliberately rather than starting with whatever is most visible. S39

> **If a candidate workflow fails two of these, it is not your beachhead.**

AMPLIFICATION

accenture

# AI amplifies the system you already have

MATURITY

**Bad at virtual machine management, then moved to cloud? Cloud let you be bad at a scale you had never seen. AI does the same to a weak workflow.**

- Security
- Reliability
- Performance
- Cost effectiveness
- Operational excellence

VERIFIABILITY

**You cannot automate a workflow whose output you cannot verify. If you cannot test at pipeline speed without AI, you cannot with it.**

Code produces testable artifacts. Data work qualifies when evaluations play that role.

> A dozen agents at ten times human speed are worth nothing if every decision chokes through one reviewer. Verification methodology is what eventually lets some decisions move without a human, and it does not exist yet for most work.

THE CORE IDEA

accenture

# Guardrails the AI enforces on itself are bypassable. Guardrails the infrastructure enforces are not

The DoD defined trustworthy in 2020 as five principles. The one that governs AI you turn loose in a mission is **Governable** : the system detects unintended behavior and can disengage on command. S6 S7

| PROMPT-LEVEL GUARDRAIL | INFRASTRUCTURE-LEVEL GUARDRAIL |
| --- | --- |
| Probabilistic: the same input can produce a different output | Deterministic: the same rule gives the same result, every time |
| The model chooses to comply | The action is blocked before the model decides |
| Jailbreakable | Enforced |
| Drifts with every model change | Logged |

> **Governable in practice means enforced by infrastructure**

A PROVEN PATTERN

accenture

# The Department of the Air Force already proved this pattern with continuous authorization

## 2018

### cATO, Air Force origin

Continuous monitoring beats periodic review

## 2025

### SWFT S8 S9

Automated, machine-readable evidence replaces the paper package

## Now

### AI behavior

Same logic, applied to AI behavior instead of software compliance

DEMONSTRATION TWO | OUTFIGHT

# Guardrails that hold

The question: how do I become certain.

accenture

Left: the request. Right: the tool calls and the policy decision.

POSTER SET > dafitc-03-coach.mp4 > video pending

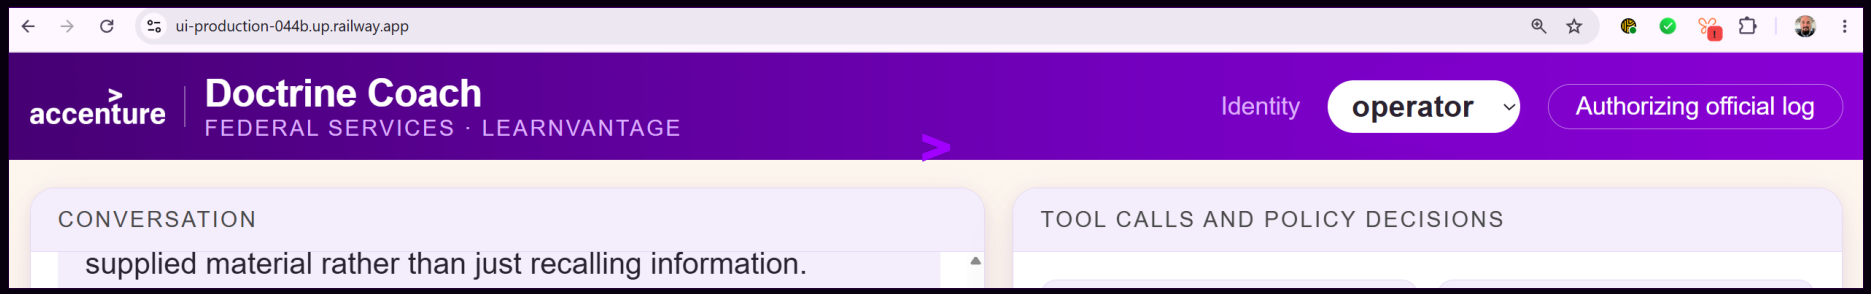


01 A request inside policy. The call executes and the answer returns.

02 **A request outside policy. The operator pushes. The model agrees to answer anyway.**

03 Denied at the policy layer before the call executes. The model complied. It did not matter.

FALLBAC
K STILLS

01 a request inside policy completes

02 the operator pushes, the model agrees

03 DENY at the policy layer

04 the log entry an authorizing official reads

accenture

# Redacted at the proxy, before the model sees it

| | | | |
|---|---|---|---|
| GUARDRAILS OFF | **Document ingested**<br>Sensitive content included | > | **Nothing is evaluated**<br>The document enters the corpus as-is | > | **Fully visible to the model**<br>Anything can surface in an answer |
| GUARDRAILS ON | **Same document ingested**<br>Nothing about the document changed | > | **Redacted at the proxy**<br>Sensitive content stripped before the model sees it | > | **Model sees the redacted version only**<br>The restricted passage never enters context |

01 A document with sensitive content is ingested. Everything in it is visible to the model.

02 Guardrails are turned on at the proxy.

03 **The same document is ingested again. The sensitive content never reaches the model.**

DEMONSTRATION THREE

# The system will not surface what you are not cleared to see

The question: how do I inject AI into my systems.

accenture

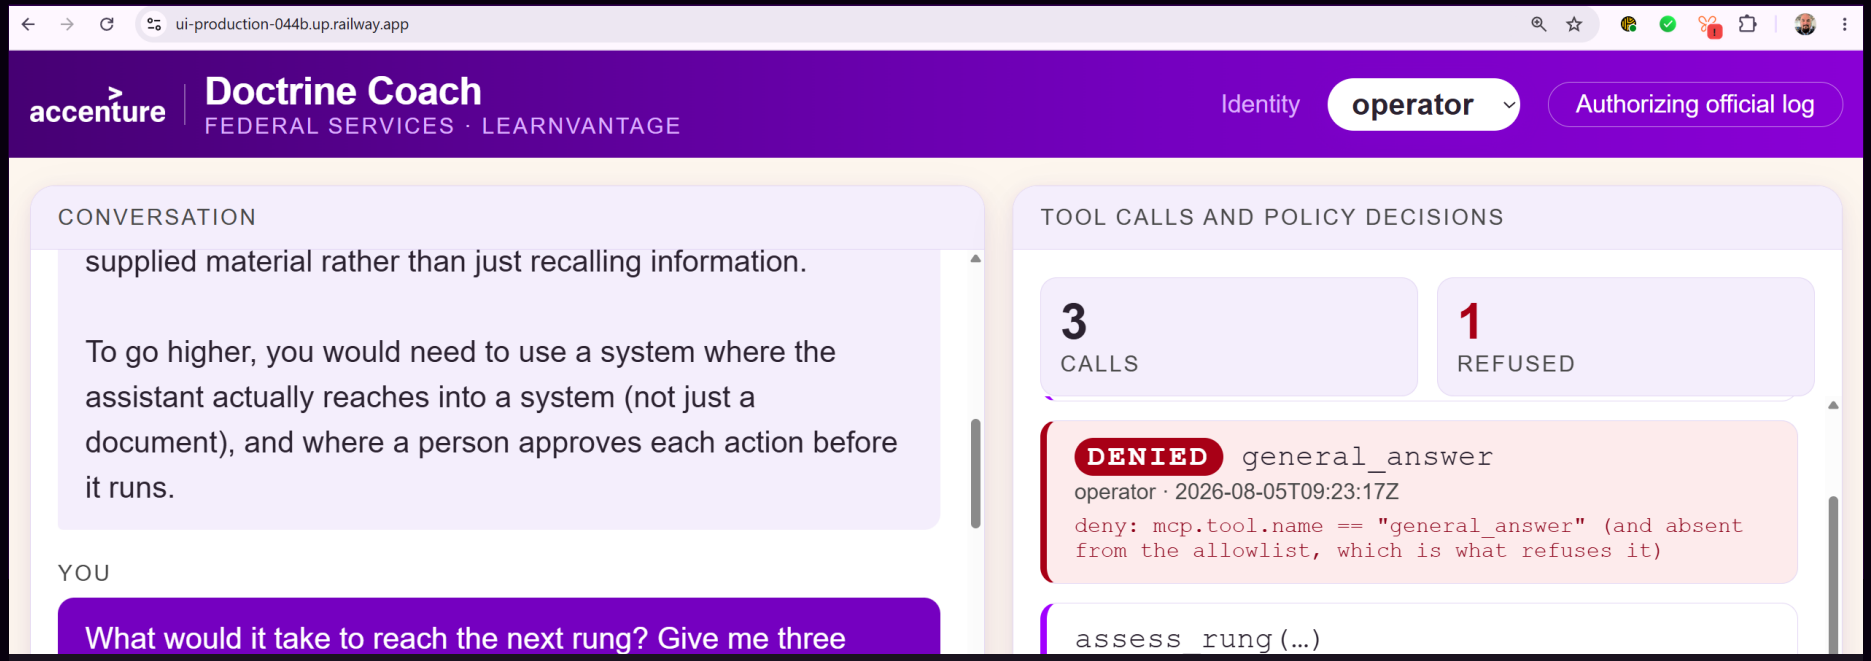


01 A document with restricted content enters the corpus. Sensitive content is redacted before ingestion, not after retrieval. S31

02 Two users ask the same question. Retrieval is filtered by role, at retrieval time.

03 **The restricted passage never enters the second user's context. The attempt is logged.** S32

BOUNDARIES

accenture

# It amplifies what is already there

- **It does not turn a novice into a specialist**
- **It does not work on doctrine you cannot sanitize and clear**
- **It does not remove the human on the mission call** S27
- **It does not verify that the advice was correct. Enforcement is not accuracy.**

**WHAT IT DOES**

Personalized learning enables skill and outcome development at the point of need

The guardrail layer deterministically stops unauthorized actions, and should be the first and last stop on what an agent can do

Guardrails depend on authorization controls that should be evaluated whether or not a model is in the loop

THE METHOD

accenture

# You work backwards from the outcome, not forwards from the tool

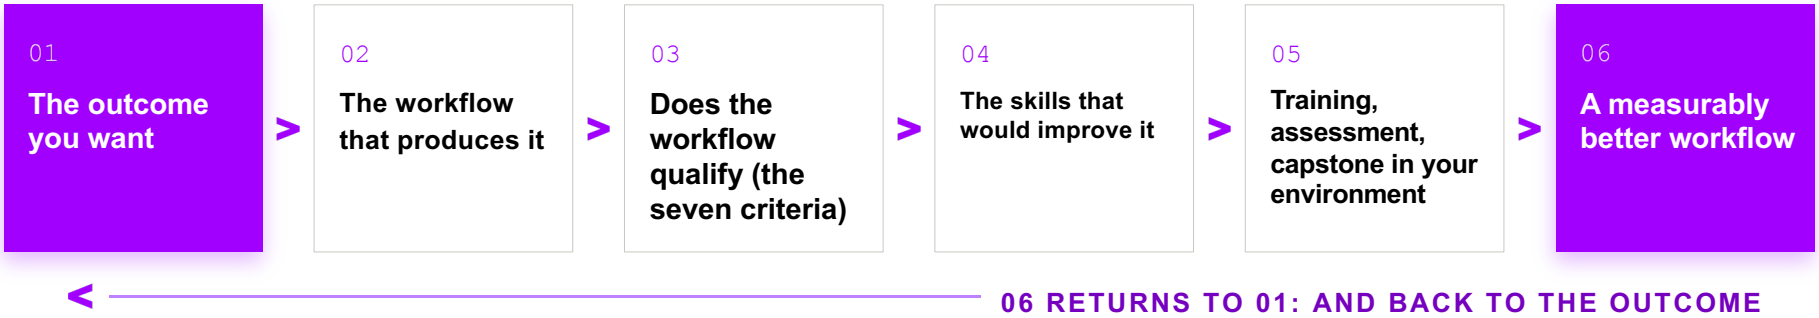


> Every step is measurable. That is what makes the outcome a measurement instead of a claim.

# Both modes serve both tracks. What changes is the tier you are building.

## LearnVantage adaptive digital training

BREADTH

Self-paced role paths, at scale

**Rungs one and two, at scale, across the whole force.**

> **Knowledge acquisition, and knowledge in the context of a specific workflow. Tier one, and part of tier two.**

MEDIA PENDING > dafitc-01-breadth.mp4 > drop-in ready

Adaptive digital training. A self-paced role path, in the product.

## LearnVantage hands-on technical skilling

DEPTH

Instructor-led, lab-based technical practice

**Rungs three and four, hands-on, where a mistake costs nothing.**

> **Deep immersion, capstones, outcome confirmation under expert evaluation. Tier two into tier three.**

MEDIA PENDING > dafitc-02-depth.mp4 > drop-in ready

Hands-on technical skilling. A live lab environment, in the product.

### TWO AUDIENCES, ONE SYSTEM

**AI-enabled operators and specialists.** Role paths grounded in mission doctrine and TTPs.

**Cyber defenders.** The same paths grounded in defensive playbooks and response procedures. The corpus changes. The enforcement layer does not.

WHAT WE JUST DID

accenture

# One workflow at a time

Demonstration one

## Personalized learning and doctrine-grounded coaching

How a person converts training into skill, and skill into outcome.

Demonstration two

## Infrastructure guardrails

The tooling you already run can deterministically constrain what an agent does.

Demonstration three

## Semantic sanitization and authorization

Beyond infrastructure guardrails: meaning-level evaluation of sensitive material, and enforcement of whether this requester can see it at all.

> **Training to skill to outcome — and the seven gates before AI enters any workflow**

accenture

# AI-Ready Warfighters. Commercial power. Mission speed.

A force that learns faster, adapts faster, and employs AI safely at the digital edge. Faster to competency. Grounded in your doctrine. Governable by design.

- **Connect with me to discuss speed-to-competency for AI-intensive missions.**
- **See it live: the coach and its guardrails, running at booth 150, closest to the entrance.**

**Accenture Federal Services delivers results that move the force forward.**

**Thank you.** Questions, now or at booth 150.

Michael Forrester / Accenture Federal Services / michael.r.forrester@accenture.com

APPENDIX

accenture

# Sources, S1 to S8


S1 **DAF launches plan to bolster AI workforce.** af.mil, April 2026

S2 **AI Ready by 2025, AI Competitive by 2027.** MeriTalk

S3 **GenAI.mil: 1.5M daily users; 50,000 waitlisted for platform onboarding.** Pentagon CTO, Hudson Institute; DefenseScoop, June 12, 2026

S4 **100,000 agents created on GenAI.mil.** DefenseScoop, April 23, 2026 https://defensescoop.com/2026/04/23/pentagon-uses-genai-mil-to-create-agents/

S5 **3.5M personnel in scope; platform launched December 2025.** Breaking Defense, February 2026

S6 **DoD adopts 5 principles of artificial intelligence ethics.** Army.mil, February 2020

S7 **Responsible AI Strategy and Implementation Pathway.** DoD, June 2022

S8 **Software Fast Track initiative, DoD CIO.** war.gov, 2025


URLs shown for sources verified August 3, 2026 · full list in the run-of-show leave-behind

APPENDIX

accenture

# Sources, S9 to S17


S9 **The DoD SWFT Initiative and the Promise of cATO Fulfilled.** Anchore

S10 **$77M DAF ICAM support task, AFLCMC.** OrangeSlices

S11 **DoD replaces paper-based access requests with automated ICAM workflow.** Federal News Network, February 2026

S12 **DAFITC 2026: Accelerating the Digital Edge, From Enterprise to Orbit.** Maxwell AFB

S13 **Accelerating the Adoption of Software and AI Agent Identity and Authorization.** NIST NCCoE concept paper, February 5, 2026

S14 **GenAI.mil model composition: Gemini for Government, OpenAI, Grok.** DefenseScoop, December 2025 through June 2026; Breaking Defense, February 2026

S17 **Randomized controlled trial on experienced developer productivity with AI tools.** METR, arXiv:2507.09089, 2025 https://arxiv.org/abs/2507.09089


URLs shown for sources verified August 3, 2026 · full list in the run-of-show leave-behind

APPENDIX

accenture

# Sources, S18 to S37


S18 **GenAI.mil at nearly 1.7M users and 100,000+ custom agents.** DoD CDAO Cameron Stanley, AWS Summit DC. Defense One, July 2, 2026
https://www.defenseone.com/technology/2026/07/genaimil-records-almost-17m-users-plans-new-model-additions/414569/

S19 **NSPM-11, Artificial Intelligence in the National Security Enterprise.** The White House, June 5, 2026
https://www.whitehouse.gov/presidential-actions/2026/06/national-security-presidential-memorandum-nspm-11/

S20 **Department of War AI strategy memoranda.** January 9, 2026

S21 **DAF AI Strategy.** Cleared for open publication April 17, 2026

S22 **DAF AI Hiring and Talent Development Plan, Total Force.** April 2026

S23 **The GenAI Divide: State of AI in Business 2025.** MIT Project NANDA, July 2025. Cross-sectional industry study

S24 **Over 40 percent of agentic AI projects will be canceled by end of 2027.** Gartner, June 25, 2025. Forecast

S25 **Artificial Intelligence Acquisitions.** GAO-26-107859, April 7, 2026
https://www.gao.gov/products/gao-26-107859

S26 **Economic limits of AI job automation in vision tasks.** MIT CSAIL, 2024
https://www.csail.mit.edu/news/rethinking-ais-impact-mit-csail-study-reveals-economic-limits-job-automation

S27 **Air Force Doctrine Note 25-1, Artificial Intelligence.** Published April 8, 2025
https://www.doctrine.af.mil/Portals/61/documents/AFDN_25-1/AFDN%2025-1%20Artificial%20Intelligence.pdf

S28 **Autonomy and Agency in Agentic AI. Two-axis autonomy and agency design space, five levels each, plus six architectural tactics.** Sahoo, Pula, Rajan, Chaudhary. arXiv 2605.12105, May 12, 2026
https://arxiv.org/abs/2605.12105

S29 **Measuring AI agent autonomy in practice.** Anthropic, 2026. Vendor telemetry

S30 **Nvidia VP: compute cost now exceeds salary spend for the same team.** Bryan Catanzaro. Axios, April 26, 2026
https://www.axios.com/2026/04/26/ai-cost-human-workers

S31 **OWASP Top 10 for LLM Applications.** 2025

S32 **NIST AI RMF, AI 100-1, and NIST SP 800-53 Rev. 5.** NIST

S33 **Inference on a 1T-parameter LLM costs over 90 percent less by 2030.** Gartner, March 25, 2026. Forecast
https://www.gartner.com/en/newsroom/press-releases/2026-03-25-gartner-predicts-that-by-2030-performing-inference-on-an-llm-with-1-trillion-parameters-will-cost-genai-providers-over-90-percent-less-than-in-2025

S35 **AI Fluency framework, version 1.1.** Dakan and Feller, 2025

S36 **Customer-facing ordering bot completes programming homework.** Media reporting, 2026

S37 **PANDA, Predictive Analytics and Decision Assistant. Rapid Sustainment Office** Selected as an Air Force system of record. DefenseScoop, May 10, 2023
https://defensescoop.com/2023/05/10/air-force-selects-ai-enabled-predictive-maintenance-program-as-system-of-record/


URLs shown for sources verified August 3, 2026 · full list in the run-of-show leave-behind