COURSE SYLLABUS

SL 4H

USAREUR-AF

COURSE SYLLABUS — SL 4H: AI ENGINEER

Maven Smart System (MSS) — USAREUR-AF

HEADQUARTERS
UNITED STATES ARMY EUROPE AND AFRICA
(USAREUR-AF)
Wiesbaden, Germany

DRAFT — NOT FOR OFFICIAL USE. FOR TRAINING PLANNING PURPOSES ONLY.

26 MARCH 2026

DRAFT — UNOFFICIAL — NOT FOR OPERATIONAL USE

COURSE SYLLABUS — SL 4H: AI ENGINEER

MAVEN SMART SYSTEM (MSS) — USAREUR-AF

Field	Detail
Level	SL 4H (AI Engineer Specialist Track)
Duration	5 days (40 hours)
Prerequisites	SL 1, SL 2, SL 3 (all Go evaluations on file — **REQUIRED**, not recommended); Data Literacy Technical Reference (required); working Python proficiency; familiarity with prompt engineering concepts recommended
Audience	AI/ML specialists, data engineers building AI-enabled pipelines on MSS
Format	Instructor-led lab + safety seminar + practical exercise
Location	MSS Training Environment (AIP Logic and Agent Studio access required)

BLUF: SL 4H teaches AI engineers to author AIP Logic workflows, configure Agent Studio agents, build RAG pipelines grounded in Ontology data, and deliver AI products meeting USAREUR-AF human-in-the-loop and OPSEC requirements. AIP Logic and Agent Studio are the most powerful — and highest-risk — tools in MSS. Day 1 begins with mandatory AI safety training. **This block cannot be skipped or rescheduled.**

LEARNING OBJECTIVES

#	Objective
1	Complete an AIP Authorization Checklist for a proposed workflow and identify prohibited use cases by category
2	Author an AIP Logic workflow with prompt engineering, conditional chain logic, and structured JSON output

#	Objective
3	Build multi-step AIP Logic chains with looping, parallel branches, and Ontology write Actions
4	Write Python transforms that extract and format Ontology data for AIP Logic context — including chunking and context window management
5	Build an LLM integration pipeline with Ontology grounding and retrieval-augmented generation (RAG)
6	Configure and test an Agent Studio agent with multiple tools, defined memory scope, and tool-use authorization controls
7	Test and red-team an AIP Logic workflow against the USAREUR-AF AI Output Validation Framework
8	Deploy an AIP Logic workflow to production with monitoring, alerting, and rollback capability

PRE-COURSE CHECKLIST

Complete **7–10+ duty days before Day 1:**

- ☐ Request AIP Logic **authoring** access from C2DAO — AIP Logic configuration access (SL 3 level) is insufficient; you need authoring privileges
- ☐ Request Agent Studio access from C2DAO
- ☐ Confirm both access levels are active and working before Day 1
- ☐ Read TM-40H, Chapter 1 (Introduction) — full read, 30 min (**required, not optional**)
- ☐ Read TM-40H, Chapter 6, Sections 6-1 and 6-2 (AI Safety) — read before the safety block, not after
- ☐ Read TM-40H, Appendix B (Prohibited AI Use Cases) — memorize the prohibited list

DAILY SCHEDULE

Day 1 — AI Safety and AIP Platform Architecture

(Safety block mandatory — no exceptions, no rescheduling)

Time	Block	Method	Content
0800–1000	1	Seminar	**AI safety:** human-in-the-loop requirements, OPSEC for AI, prohibited use cases by category, Army CIO policy (April 2024), what happens when AI outputs are wrong

Time	Block	Method	Content
1000–1015	—	Break	
1015–1200	2	Lecture	AIP platform architecture: Logic, Agent Studio, Code Workspaces, LLM endpoints — the full stack and how components connect
1200–1300	—	Lunch	
1300–1500	3	Lab	AIP Logic: authoring first workflow — prompt template, input configuration, output binding, running a test
1500–1515	—	Break	
1515–1700	4	Lab	AIP Logic: conditional chains, error handling, structured JSON output

Evening reading: TM-40H, Chapter 3 (AIP Logic workflow authoring); Chapter 6 Sections 6-3/6-4 (OPSEC for AI outputs).

Day 2 — Advanced AIP Logic and Python Transforms

Time	Block	Method	Content
0800–0830	—	Review	Day 1 questions; the non-negotiable: human-in-the-loop checkpoint design
0830–1030	5	Lab	AIP Logic: multi-step chains, looping, parallel branches, Action integration
1030–1045	—	Break	
1045–1200	6	Lab	Python transforms for AIP: extracting Ontology data, context formatting for operational terminology (DTG, DODAAC, MTOE — LLMs do not know these by default)
1200–1300	—	Lunch	
1300–1500	7	Lab	Context management: chunking strategies, context window limits, handling large datasets without exceeding LLM context
1500–1515	—	Break	

Time	Block	Method	Content
1515–1700	8	Lab	Ontology integration: writing AIP Logic outputs to Object properties via Actions; implementing a human review queue for uncertain outputs

Evening reading: TM-40H, Chapter 5 (LLM Integration Patterns — RAG architecture section).

Day 3 — RAG Architecture and LLM Integration

Time	Block	Method	Content
0800–0830	—	Review	Day 2 questions; review queue design patterns — confirm trainees can implement the Draft→Review→Publish workflow
0830–1030	9	Lab	RAG architecture: semantic search setup, retrieval from Ontology Objects, context construction for grounded responses
1030–1045	—	Break	
1045–1200	10	Lab	RAG pipeline build: structured retrieval → context formatting → LLM prompt → structured JSON output → Ontology write with review gate
1200–1300	—	Lunch	
1300–1500	11	Lab	RAG quality: evaluating retrieval relevance, grounding failures, hallucination detection patterns for operational data
1500–1515	—	Break	
1515–1700	12	Lab	End-to-end workflow practice: build a complete AIP Logic workflow with RAG, human review queue, and Ontology write — instructor coaching

Evening reading: TM-40H, Appendix A (AIP Authorization Checklist) — read in full; you will complete this on Day 4/5.

Day 4 — Agent Studio and Evaluation

Time	Block	Method	Content
0800–0830	—	Review	Day 3 questions; RAG grounding failure patterns debrief

Time	Block	Method	Content
0830–1100	13	Lab	Agent Studio: agent architecture, tool definition, tool-use authorization controls, memory scope configuration, orchestration patterns
1100–1115	—	Break	
1115–1200	14	Lab	Agent Studio: testing tool use; confirming authorization controls prevent out-of-scope actions
1200–1300	—	Lunch	
1300–1500	15	Lab	Testing AI outputs: evaluation frameworks, red-teaming methodologies, adversarial prompt testing against operational data contexts
1500–1515	—	Break	
1515–1700	16	Lab	Applying the AI Output Validation Framework (TM-40H, Appendix C); completing the AIP Authorization Checklist for a practice workflow

Evening reading: Review your Day 3 RAG workflow — document 3 failure scenarios and mitigations before Day 5.

Day 5 — Deployment and Practical Exercise

Time	Block	Method	Content
0800–0900	17	Lab	Production deployment: pipeline scheduling, monitoring setup, build failure alerting, rollback procedures
0900–1000	18	Brief	Practical exercise scenario brief; planning and workflow design time (document architecture before building)
1000–1015	—	Break	
1015–1100	19	Brief	Authorization Checklist completion guidance; evaluation criteria review
1100–1200	—	Buffer	Questions / environment check
1200–1300	—	Lunch	

Time	Block	Method	Content
1300–1700	20	**Eval**	**Practical exercise:** author → test → authorize → deploy an AIP Logic workflow end-to-end

REQUIRED READING SUMMARY

When	Reading
Before Day 1 (**required**)	TM-40H, Ch 1 (Introduction)
Before Day 1 (**required**)	TM-40H, Ch 6 Sec 6-1/6-2 (AI Safety)
Before Day 1 (**required**)	TM-40H, Appendix B (Prohibited Use Cases)
Day 1 evening	TM-40H, Ch 3 (AIP Logic workflow authoring)
Day 1 evening	TM-40H, Ch 6 Sec 6-3/6-4 (OPSEC for AI)
Day 2 evening	TM-40H, Ch 5 (LLM Integration Patterns — RAG)
Day 3 evening	TM-40H, Appendix A (Authorization Checklist)
Day 4 evening	Review Day 3 workflow; document 3 failure scenarios

PRACTICAL EXERCISE

Scenario: The S2 shop receives unstructured SITREP reports in plain text. Build an AIP Logic workflow that extracts structured data from each submission, maps it to a SITREP Object Type via an Action, and routes uncertain extractions to a human review queue.

#	Task
1	Design the workflow architecture — document before building; include the human review gate in the design
2	Author the AIP Logic workflow with prompt engineering and structured JSON output; define explicit context for military terminology
3	Implement a human review queue: extractions below a confidence threshold route to a review queue Workshop application, not directly to production Objects
4	Configure the Ontology write Action with a confirmation checkpoint; confirm no write executes without human review

#	Task
5	Test against 5 provided SITREP samples — validate extraction accuracy; confirm the review queue catches uncertain outputs
6	Complete the AIP Authorization Checklist for the workflow
7	Configure monitoring and failure alerting for production deployment

Go standard: Pass 6 of 7 tasks. Evaluator confirms no write Action executes without a human checkpoint — the evaluator will attempt to trigger a write bypassing the checkpoint. Authorization Checklist complete and honest.

GO CRITERIA

Requirement	Hard Standard
Human-in-the-loop	Any workflow that writes to production Ontology Objects without a tested human checkpoint fails — regardless of everything else. The evaluator will specifically attempt to bypass the checkpoint; design the workflow so bypass is impossible by construction
Authorization Checklist	Must be completed honestly — misrepresenting the workflow fails. If the workflow touches data in a prohibited or restricted category, flag it

KEY TIPS

Risk	Guidance
AI safety block	Not a compliance checkbox — the most important content in the course. Engineers who treat it as a formality build workflows that fail in production in ways that matter. Read Appendix B before Day 1
Prompt engineering on operational data	Military terminology (DTG, DODAAC, MTOE) is not in most LLMs' training data — prompts must explicitly define these terms. Test against deliberately ambiguous inputs before the evaluation
Review queue	When in doubt, route to human review. A KPI of "zero items in the review queue" may mean uncertain extractions flow directly to production data
RAG context construction	You are not dumping an entire dataset into the context window — retrieve the most relevant objects and present them in a structured format. Read Chapter 5 grounding

Risk	Guidance
	patterns before Day 3
Authorization Checklist	Final question: "What happens if this workflow produces incorrect output?" If the answer is "nothing bad," you have not thought hard enough about downstream use

CONTINUATION

Graduates who move into enterprise AI architecture or multi-agent system design roles may pursue **SL 5H (Advanced AI Engineering)**. SL 5H covers enterprise-scale RAG infrastructure, multi-agent orchestration, AI governance frameworks, and evaluated system architecture. Prerequisites: SL 4H Go evaluation on file; 12+ months active AI engineering experience on MSS or equivalent.

ASSOCIATED EXERCISES AND ASSESSMENTS

Item	Reference
Pre-course exam	EXAM_TM40H_PRE
Post-course exam	EXAM_TM40H_POST
Practical exercise	EX_40H (EXERCISE.md + ENVIRONMENT_SETUP.md)

RELATIONSHIP TO WFF TRACKS

WFF track analysts (SL 4A through SL 4F) are the operational consumers of AI-enabled tools built in this course. AI engineers should understand the operational workflows each WFF community uses on MSS: intelligence analysts (SL 4A) consume structured extraction products and automated reporting; mission command analysts (SL 4F) use AIP Logic workflows for information synthesis and course of action support; all WFF communities benefit from human-review-gated AI products that meet USAREUR-AF safety standards.

USAREUR-AF Operational Data Team Syllabus SL 4H | Version 2.0 | March 2026