COURSE SYLLABUS

SL 5H

USAREUR-AF

COURSE SYLLABUS — SL 5H: ADVANCED AI ENGINEERING

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 5H: ADVANCED AI ENGINEERING

MAVEN SMART SYSTEM (MSS) — USAREUR-AF

Field	Detail
Level	SL 5H — Advanced AI Engineer Specialist Track
Duration	5 days (40 hours)
Prerequisites	SL 4H complete (Go evaluation on file); 12+ months active AI engineering experience on MSS or equivalent; demonstrated proficiency with Python, LLM APIs, and vector retrieval systems; SL 4M (ML Engineer) recommended
Audience	Senior AI engineers, AI architects, AI capability leads assigned to enterprise-level AI development on MSS
Format	Advanced lab + architecture review + evaluated system design
Location	MSS Training Environment (AIP Logic, Agent Studio, Python Transforms enabled; advanced model access required)

PREREQUISITE WARNING: SL 5H is not required for the majority of AI engineer billets. It is intended for personnel designing AI system architectures, retrieval infrastructure, or agent orchestration frameworks at enterprise scale. Consult your supervisor or C2DAO before enrolling.

BLUF: SL 5H addresses the architectural and governance challenges that arise when AI systems move from individual workflows to enterprise deployments — multi-agent orchestration, production RAG infrastructure, AI governance frameworks, and responsible deployment at scale. These techniques are for engineers building the systems that other analysts and staff will use.

LEARNING OBJECTIVES

#	Objective
1	Design multi-agent orchestration architectures: agent routing, memory sharing, tool hand-off, and failure recovery across agent networks
2	Build and optimize enterprise-scale RAG pipelines: chunking strategies, embedding model selection, hybrid retrieval (dense + sparse), re-ranking, and retrieval quality evaluation
3	Implement AI governance frameworks: human-in-the-loop gate design, output validation pipelines, audit logging for AI decisions, and rollback procedures
4	Design prompt engineering standards for production systems: few-shot templating, chain-of-thought patterns, output format enforcement, and adversarial robustness
5	Build AI evaluation harnesses: automated regression testing for AI outputs, ground truth dataset design, inter-rater reliability for human evaluation
6	Apply OPSEC and classification handling requirements to AI systems that ingest or generate operationally sensitive content
7	Document an AI system architecture for a technical audience: data flow, failure modes, human review gates, retraining triggers, and deprecation criteria

PRE-COURSE CHECKLIST

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

- ☐ Confirm AIP Logic advanced access (multi-agent features) with C2DAO — standard AIP Logic access is insufficient for Day 3 content
- ☐ Read TM-50H, Chapter 1 (Introduction and Scope) and Chapter 8 (Enterprise AI Architecture and Governance) before Day 1
- ☐ Prepare a 1-page description of an AI system or workflow you currently maintain — you will use this for the architecture review on Day 4

DAILY SCHEDULE

Day 1 — Enterprise RAG Architecture

Time	Block	Method	Content
0800–0900	1	Seminar	SL 5H scope and standards; enterprise vs. prototype AI systems; the gap between "demo" and "production"
0900–1100	2	Lab	RAG pipeline architecture: document ingestion, chunking strategy tradeoffs (fixed vs. semantic), metadata schema design
1100–1115	—	Break	
1115–1200	3	Lab	Embedding model selection: criteria, trade-offs, OPSEC implications of external embedding APIs vs. on-premises models
1200–1300	—	Lunch	
1300–1500	4	Lab	Hybrid retrieval: combining dense (vector) and sparse (BM25) retrieval; re-ranking pipelines; retrieval quality metrics (MRR, NDCG@k)
1500–1515	—	Break	
1515–1700	5	Lab	Retrieval evaluation: building a ground truth query set; automated retrieval quality harness

Evening reading: TM-50H, Chapter 4 (Advanced RAG Architecture).

Day 2 — Prompt Engineering for Production

Time	Block	Method	Content
0800–0830	—	Review	RAG quality issues; common retrieval failures
0830–1030	6	Lab	Prompt engineering standards: system prompt architecture, few-shot template design, output format enforcement (JSON schema, structured extraction)
1030–1045	—	Break	

Time	Block	Method	Content
1045–1200	7	Lab	Chain-of-thought patterns: when CoT improves reliability and when it does not; step-back prompting; self-consistency
1200–1300	—	Lunch	
1300–1500	8	Lab	Adversarial robustness: prompt injection, jailbreak patterns, and mitigation strategies for operational AI systems
1500–1515	—	Break	
1515–1700	9	Lab	Prompt version control and regression testing: tracking prompt changes, automated comparison of output distributions

Evening reading: TM-50H, Chapter 2 (Multi-Agent Orchestration Systems) — agent routing and failure recovery sections.

Day 3 — Multi-Agent Orchestration

Time	Block	Method	Content
0800–0830	—	Review	Prompt engineering review; common adversarial patterns seen in operational context
0830–1030	10	Lab	Multi-agent architecture: orchestrator/worker patterns, agent routing logic, capability registration
1030–1045	—	Break	
1045–1200	11	Lab	Memory architectures: short-term (context window management), long-term (vector store, structured memory), shared memory across agents
1200–1300	—	Lunch	
1300–1500	12	Lab	Failure recovery: agent timeout handling, fallback chains, dead letter queues for failed agent tasks
1500–1515	—	Break	
1515–1700	13	Lab	Tool hand-off patterns: structured tool output schemas, downstream tool input validation, circular dependency detection

Evening reading: TM-50H, Chapter 8 (Enterprise AI Architecture and Governance) — human-in-the-loop gate design.

Day 4 — AI Governance and System Design Review

Time	Block	Method	Content
0800–0900	14	Seminar	AI governance framework: human review gate design principles, audit log requirements, output validation pipeline architecture
0900–1100	15	Workshop	Participant architecture review: present your current AI system (from pre-course prep); peer and instructor critique
1100–1115	—	Break	
1115–1200	16	Workshop	Evaluating AI outputs at scale: automated regression testing design, human evaluation calibration, inter-rater reliability
1200–1300	—	Lunch	
1300–1500	17	Lab	OPSEC for AI systems: classification handling in RAG pipelines, access controls on knowledge stores, output filtering for export
1500–1515	—	Break	
1515–1700	18	Lab	AI system documentation standards: architecture diagram, failure modes register, retraining triggers, deprecation criteria

Evening reading: Prepare evaluation system design (Day 5 evaluated exercise).

Day 5 — Evaluated System Design

Time	Block	Method	Content
0800–0900	19	Brief	Evaluation scenario brief and design requirements
0900–1700	20	**Eval**	**Evaluated exercise:** Design and build a multi-agent retrieval system with governance gates; produce a system architecture document and evaluation harness

Go standard: System functional with at least 2 agents + retrieval tool; governance gate present; architecture document complete with failure modes register.

PEER ADVANCED TRACKS

Track	Relevance to SL 5H
SL 5M (Advanced ML Eng)	ML model serving integration with AI pipelines; retraining triggers from AI output quality monitoring
SL 5G (Advanced ORSA)	AI-assisted analytical products; grounding AI reasoning in quantitative ORSA outputs
SL 5L (Advanced SWE)	OSDK-backed AI system integrations; platform security requirements for AI applications consuming sensitive Ontology data

USAREUR-AF Operational Data Team Syllabus SL 5H | Version 1.0 | March 2026