PUBLICATION

EXAM-TM50L-PRE

PRE-TEST — SL 5L: ADVANCED SOFTWARE ENGINEER

Maven Smart System (MSS) — USAREUR-AF

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

26 MARCH 2026

DRAFT — UNOFFICIAL — NOT FOR OPERATIONAL USE

PRE-TEST — SL 5L: ADVANCED SOFTWARE ENGINEER

MAVEN SMART SYSTEM (MSS) — USAREUR-AF

Field	Detail
Course	SL 5L: Advanced Software Engineer
Form	Pre-Test
Level	SL 5L (Advanced Specialist)
Audience	Senior SWEs / platform architects / tech leads; prerequisite: SL 4L + production MSS SWE experience
Time Allowed	30 minutes
Passing Score	N/A — diagnostic only

INSTRUCTIONS

This diagnostic assessment establishes your baseline knowledge before training. Your score does not affect course eligibility. Answer honestly — results help the instructor tailor instruction to gaps.

SECTION 1 — MULTIPLE CHOICE

Circle the letter of the best answer. (2 points each)

1. In a multi-tenant platform architecture, "tenant isolation" ensures:

A. Each tenant has its own dedicated compute cluster B. Tenants are physically separated on different network segments C. One tenant's data is inaccessible to other tenants' applications and queries, preventing cross-tenant data exposure regardless of how queries are constructed D. Tenant administrators cannot see the platform-level configuration

2. "Content-Based Access Control" (CBAC) in a Foundry multi-tenant environment grants access based on:

A. The user's organizational role defined in Active Directory B. The tenant's subscription tier and licensed feature set C. Attributes of the data content itself (classification, unit, sensitivity markings) rather than only user role or resource location D. The time of day and network location of the access request

3. A "bulk operation" in the Platform SDK differs from per-record operations because:

A. Bulk operations bypass access control checks for performance B. Bulk operations write directly to the database, bypassing Ontology validation C. Bulk operations require elevated permissions not available to standard SDK clients D. Bulk operations process multiple records in a single API call, reducing network overhead and improving throughput compared to N individual calls

4. "Static Application Security Testing" (SAST) in a DevSecOps pipeline:

A. Tests the running application by simulating user interactions B. Analyzes source code for known vulnerability patterns without executing the code C. Scans container images for outdated dependencies with known CVEs D. Validates that the application handles malformed HTTP requests correctly

5. "Dynamic Application Security Testing" (DAST) in a DevSecOps pipeline:

A. Analyzes source code for known vulnerability patterns B. Validates API schemas against a defined specification C. Scans infrastructure-as-code templates for misconfigurations D. Tests the running application by sending malformed or adversarial inputs to discover runtime vulnerabilities

6. In event streaming architecture, a Kafka "consumer group" processes messages by:

A. Distributing topic partitions across consumers in the group so that each message is processed by exactly one consumer — enabling parallel processing at scale B. Each consumer in the group independently reading all messages from the topic C. Having a single designated consumer read messages and broadcast them to others in the group D. Consuming only the most recent message on each partition, discarding older messages

7. gRPC is preferred over REST for high-throughput inter-service communication because:

A. gRPC uses HTTP/2 with binary Protocol Buffer serialization, providing lower latency and higher throughput than REST/JSON over HTTP/1.1 B. gRPC provides built-in authentication that REST does not support C. gRPC is easier to debug than REST due to its binary format D. gRPC automatically handles retry logic and circuit breaking

8. In an Ontology-level CI pipeline for Foundry, "Ontology CI" validates:

A. That all pipeline code compiles without errors B. That the Ontology schema conforms to USAREUR-AF naming conventions C. That all new features are covered by unit tests before merge D. That changes to Object Types, Link Types, and properties do not break existing downstream consumers, queries, or Action validators

9. An "ATO (Authority to Operate)" package for a Foundry-based system typically requires:

A. A penetration test report and a signed memorandum from the unit commander B. A system description, data flow diagrams, security controls mapping (NIST SP 800-53 or equivalent), network boundary diagram, and risk assessment C. Vendor certification that the platform meets DoD IL-4 or IL-5 requirements D. Annual revalidation by the MSS program office

10. "Branch automation" in the Platform SDK refers to:

A. Automatically merging branches after a defined time period without human review B. Automatically creating feature branches when a Jira ticket is created C. Programmatic creation, management, and merging of Foundry dataset branches via SDK calls, enabling automated data management workflows D. Running pipeline builds on branches without requiring a data steward to trigger them manually

11. "Cache invalidation" is considered one of the hardest problems in software engineering because:

A. Caches require complex cryptographic operations to maintain consistency B. Determining when cached data is stale — and ensuring all consumers see fresh data without unnecessary cache misses — requires careful coordination between write and read paths C. Cache size limits require constant tuning as data volumes change D. Different programming languages implement cache invalidation incompatibly

12. An OWASP-informed secure code review of a Foundry OSDK integration would specifically check for which vulnerability category?

A. Injection vulnerabilities — specifically unsanitized user input being incorporated into Ontology queries or Action parameters that could alter query logic B. Buffer overflow vulnerabilities in TypeScript string handling C. Cross-site scripting (XSS) in server-rendered HTML templates D. Broken cryptographic algorithm usage in data-at-rest encryption

13. In a high-throughput ingestion architecture where an external Army system sends 50,000 records per minute to a Foundry dataset, the ingestion pattern should:

A. Write each record individually as it arrives to minimize latency B. Buffer incoming records into batches and write in bulk transactions to reduce API call overhead and maintain dataset transaction integrity C. Require a human to review each batch before committing to the dataset D. Store records in a local queue and process them once per hour to avoid overloading the platform

14. A "write-through cache" strategy means:

A. Writes go to the cache only; the database is updated asynchronously B. Reads populate the cache on the first miss; subsequent reads serve from cache C. Writes update both the cache and the underlying database synchronously — the cache is always consistent with the database D. Writes bypass the cache and go directly to the database; the cache is invalidated on write

15. Platform governance at the SL 5L level includes which responsibility that is NOT present at SL 4L?

A. Setting and enforcing platform-wide coding standards, conducting architecture reviews of other engineers' designs, and onboarding and developing junior engineers to SL 4L competency B. Writing TypeScript code that passes Ontology CI checks C. Submitting pull requests that pass the standard code review checklist D. Ensuring personal work is covered by the minimum required test cases

---

SECTION 2 — SHORT ANSWER

Answer in 2–5 sentences. (6 points each)

SA-1. Describe the difference between SAST and DAST in a DevSecOps pipeline. For each, give one specific vulnerability that it would and would not catch in a Foundry OSDK integration.

SA-2. Explain why gRPC is appropriate for a high-throughput sensor data ingestion pipeline feeding the MSS Ontology, but REST would be appropriate for the Slate application that allows G3 staff to query vehicle positions. Describe the key technical trade-offs driving each choice.

SA-3. Describe the multi-tenant isolation requirements for a USAREUR-AF MSS deployment where division-level data must be accessible within each division but not visible cross-division. What CBAC configuration would you design, and what AR 25-2 compliance considerations apply?

SA-4. A junior SWE on your team proposes bypassing CBAC in the staging environment "just for testing" to avoid setting up test data for each tenant. As the SL 5L platform lead, describe your response and explain why this is never acceptable even in non-production environments.

SA-5. You are onboarding a new SL 4L qualified engineer to the USAREUR-AF MSS platform. Describe the four most important platform-specific concepts you would cover in their onboarding, and explain why each is critical to safe, effective work on the platform.

SCORING SUMMARY

Section	Questions	Points Each	Total Points
Multiple Choice	15	2	30
Short Answer	5	6	30
Total	—	—	**60**

Passing: N/A — Pre-test is diagnostic only.

ANSWER KEY — INSTRUCTOR USE ONLY

Do not distribute to students.

Multiple Choice: 1. C — Tenant isolation = data inaccessibility across tenants regardless of query construction. 2. C — CBAC grants access based on content attributes (classification, unit, sensitivity). 3. D — Bulk operations process multiple records in a single API call, reducing overhead vs. N individual calls. 4. B — SAST analyzes source code without executing it. 5. D — DAST tests the running application with adversarial inputs. 6. A — Consumer groups distribute partitions so each message is processed by exactly one consumer. 7. A — gRPC uses HTTP/2 + binary Protocol Buffers: lower latency and higher throughput than REST/JSON. 8. D — Ontology CI validates that schema changes do not break existing downstream consumers. 9. B — ATO requires system description, data flows, security controls mapping, boundary diagram, risk assessment. 10. C — Branch automation = programmatic SDK-driven branch

creation, management, and merging. 11. B — Cache invalidation is hard because determining when data is stale and coordinating read/write paths requires careful design. 12. A — Injection vulnerabilities (unsanitized input in Ontology queries/Action parameters) are the primary OWASP concern for OSDK integrations. 13. B — Batch writes reduce API call overhead and maintain transaction integrity. 14. C — Write-through: writes update both cache and database synchronously; cache always consistent. 15. A — Platform governance = setting standards, architecture reviews, onboarding/developing junior engineers — not present at SL 4L.

Short Answer Guidance:

SA-1. Full credit: SAST — analyzes code without running it; catches: hardcoded credentials in TypeScript source, unsanitized input passed to Ontology query; would NOT catch: runtime injection from externally-supplied data that only exists at execution time; DAST — tests running application with adversarial inputs; catches: runtime injection vulnerabilities that only manifest with malformed input, authentication bypass under specific runtime conditions; would NOT catch: insecure code patterns that never produce a runtime error (e.g., unnecessarily broad permissions defined in code). Both tools with one catch and one miss each, in Foundry/OSDK context, required for full credit.

SA-2. Full credit: gRPC for ingestion — high message volume (50K+ records/min) requires binary serialization and HTTP/2 multiplexing to minimize per-message overhead; synchronous REST/JSON would produce unacceptable latency and throughput under this load; trade-offs: gRPC requires Protocol Buffer schema maintenance and is harder to inspect/debug; REST for G3 Slate query — broad client compatibility (browser-based Slate uses standard HTTP), moderate request volume, human-initiated queries tolerate higher per-request latency; REST's self-describing JSON responses are easier for the application developer to work with; trade-offs: REST/JSON has higher per-request overhead than gRPC but is appropriate at G3 query volumes. Both choices must include the specific technical trade-off driving the decision.

SA-3. Full credit: CBAC configuration — define a content attribute (e.g., `owning_division`) on all Object Types containing division-specific data; configure CBAC policies so each division's data is accessible only to users with a matching `division` attribute in their profile; cross-division access requires an explicit elevated grant from the data steward; AR 25-2 compliance — Army Regulation 25-2 (Information Assurance) requires that access controls on Army IT systems be implemented to prevent unauthorized access; CBAC configuration must be documented, reviewed, and approved as part of the system's ATO; any CBAC policy change requires C2DAO sign-off and a new or updated security review. Must include CBAC design, the cross-division access control, and AR 25-2 reference.

SA-4. Full credit: response — reject the proposal immediately and explain: CBAC bypass is NEVER acceptable in any environment (not staging, not testing, not under operational pressure); reasons: (1) staging environments often use production-equivalent data or data with the same classification boundaries; bypassing CBAC in staging normalizes the bypass pattern and creates risk of the same approach being attempted in production; (2) engineers who develop on a CBAC-bypassed environment may not test the actual CBAC-controlled behavior — tests pass in staging but fail in production; (3) AR

25-2 compliance does not have a "staging exception" — the access control framework applies to all environments handling Army data; correct alternative: set up a test data fixture for each tenant in the staging environment. Full credit requires all three reasons AND the correct alternative.

SA-5. Full credit: any four from — (1) CBAC configuration and tenant isolation — how the platform enforces data isolation between units; why critical: misconfigured CBAC is a silent data exposure risk; (2) C2DAO branch and promotion workflow — all Ontology changes require a data steward review and sign-off; why critical: direct production changes are not possible; (3) Platform SDK credential handling — SDK credentials are high-privilege operational secrets; exposure requires immediate reporting and rotation; why critical: a leaked credential has theater-wide data access; (4) ResourceIterator pagination requirement — all OSDK queries must consume all pages; why critical: partial consumption silently returns incomplete datasets; (5) Ontology CI pipeline — every schema change runs automated CI; why critical: breaking schema changes affect all downstream consumers; (6) N+1 query anti-pattern at platform scale — per-object queries degrade platform performance for all users. Four items with specific "why critical" explanation each required.

USAREUR-AF Operational Data Team TM-50L Pre-Test | Version 1.0 | March 2026