PUBLICATION

EXAM-TM50L-POST

POST-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

POST-TEST — SL 5L: ADVANCED SOFTWARE ENGINEER

MAVEN SMART SYSTEM (MSS) — USAREUR-AF

Field	Detail
Course	SL 5L: Advanced Software Engineer
Form	Post-Test
Level	SL 5L (Advanced Specialist)
Audience	Senior SWEs / platform architects / tech leads; prerequisite: SL 4L + production MSS SWE experience
Time Allowed	45 minutes
Passing Score	70% (42/60)

INSTRUCTIONS

This assessment evaluates mastery of course learning objectives. A passing score of 70% is required to receive credit. Complete independently without reference to training materials.

SECTION 1 — MULTIPLE CHOICE

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

1. A junior engineer on your team proposes bypassing CBAC on a staging tenant "just until the sprint demo is done" because the test data setup is taking too long. Per SL 5L, your response is:

A. Allow it for the staging environment only — CBAC bypass is acceptable in non-production tenants B. Allow it if the data steward approves it in writing for the duration of the sprint C. Reject the bypass without exception — CBAC bypass is NEVER acceptable in any environment; the correct fix is to set up properly scoped test data for each tenant D. Allow it with a formal risk acceptance signed by the tech lead

2. The Platform SDK credentials used for an automated ingestion pipeline are discovered in a commit to the team's Git repository. Per SL 5L, the IMMEDIATE required actions are:

A. Delete the commit and rebase history to remove the credentials; notify the team to pull the latest branch B. Report the exposure as a security incident immediately, rotate the credentials, audit the access log for any unauthorized use during the exposure window, and investigate how the credentials reached the repository C. Reset the credentials and push a new commit that uses environment variables instead D. Assess the repository's access permissions — if only cleared team members have access, no further action is required

3. An Ontology CI pipeline must validate which of the following before allowing an Object Type schema change to merge?

A. That the change has been approved by the program manager B. That no existing downstream OSDK queries, Action validators, TypeScript FOOs, or Pipeline Builder pipelines will break due to the schema change — specifically checking for removed or renamed properties that are referenced downstream C. That the change was authored by a SL 5L qualified engineer D. That the change passes SAST scanning for security vulnerabilities

4. Per SL 5L, a C2DAO sign-off is required for which of the following operations?

A. Adding a new Widget to a Workshop application B. Deploying a new Contour analysis to the production project C. Any Pipeline Builder pipeline change that affects a curated dataset D. Both Ontology branch merges to production AND changes to CBAC policy configurations

5. A high-throughput ingestion pipeline consuming events from a Kinesis stream into a Foundry dataset should implement which pattern to prevent partial writes from corrupting downstream consumers?

A. Buffer events into batches, write the batch in a single Foundry write transaction (atomic commit), and only advance the Kinesis checkpoint after the transaction commits successfully B. Write each event individually as it is consumed from the stream C. Write events directly to the Ontology Object Type to skip the intermediate dataset step D. Use a separate pipeline to validate each event before it is written to the dataset

6. You are profiling a Foundry OSDK query that returns all Vehicle Objects for V Corps (approximately 18,000 records). The query takes 47 seconds. The FIRST profiling step is:

A. Add an index to the `unit_designation` property on the Vehicle Object Type B. Reduce the number of properties returned per object to a minimal subset C. Identify whether the bottleneck is in the OSDK client pagination overhead, the Ontology query execution, or the network transfer — by measuring latency at each stage separately before applying any optimization D. Split the query into parallel requests by unit echelon to reduce per-query record counts

7. A "cache-aside" caching strategy for frequently queried Ontology data means:

A. The application writes to the cache and database simultaneously on every write B. The cache is pre-populated at application startup with all expected queries C. On a cache miss, the application fetches from the Ontology, populates the cache, and serves the result; on cache hit, the application serves from cache without querying the Ontology D. Writes invalidate the cache asynchronously after the database write completes

8. An external Army logistics system will push real-time supply requisition data via gRPC to a USAREUR-AF MSS ingestion endpoint. Before the external feed is activated in production, which step is required per SL 5L?

A. A load test demonstrating the endpoint can handle peak message volume B. A data steward review confirming the Ontology schema can accommodate the new data fields C. A GO-level authorization memorandum approving the external integration D. A security review of the external data feed — validating data provenance, input schema, and ensuring that untrusted external data is validated and sanitized before entering the Foundry Ontology

9. Performance profiling of a production Foundry environment per SL 5L requires:

A. Running profiling queries directly against the production Ontology without restriction — performance data is non-sensitive B. Using only pre-approved profiling queries from the USAREUR-AF approved query library C. Coordination with the C2DAO data steward before running profiling queries in production — resource-intensive profiling can degrade platform performance for operational users D. Conducting all profiling during off-peak hours without additional coordination

10. A SL 5L security assessment of a Slate application reveals that a user-supplied unit designation string is interpolated directly into an Ontology Object filter query without sanitization. This is classified as:

A. An informational finding — TypeScript type safety prevents this from being exploitable B. A critical injection vulnerability — unsanitized user input in query logic can alter query behavior, potentially exposing records outside the user's authorized scope C. A medium-severity finding — Ontology queries are parameterized by the OSDK and not directly injectable D. A low-severity finding — the application is only accessible to cleared users who would not attempt injection

11. The SL 5L peer review requirement for all platform-level work mandates that the reviewer must be:

A. Any engineer on the team who did not author the code B. A second SL 5L qualified engineer — the complexity of platform-level CBAC configurations, Ontology CI changes, and SDK integrations requires reviewer expertise at the same qualification level C. The tech lead, regardless of their TM qualification level D. An engineer from outside the team to prevent conflict of interest

12. In a multi-tenant USAREUR-AF MSS deployment, cross-tenant data bleed would violate which Army Regulation?

A. AR 25-2 (Army Cybersecurity), which requires that information systems protect data from unauthorized access — cross-tenant bleed constitutes unauthorized access to another unit's operational data B. AR 600-8-2 (Suspension of Favorable Personnel Actions) C. AR 380-5 (Department of the Army Information Security Program) D. AR 25-400-2 (The Army Records Information Management System)

13. An Ontology branch automation workflow that creates a new dataset branch, runs pipeline transforms, validates output quality, and merges to main without human review is:

A. Non-compliant with C2DAO requirements — branch merges to production require C2DAO sign-off as a human-in-the-loop step regardless of automated check results B. Acceptable if all automated quality checks pass — this is the purpose of branch automation C. Acceptable for dataset branches; non-compliant only for Ontology schema branches D. Acceptable if the pipeline is designated as low-risk by the data steward

14. When onboarding a new SL 4L engineer to the platform, the FIRST topic to cover per SL 5L onboarding standards is:

A. The OSDK authentication pattern and ResourceIterator pagination B. The Ontology CI pipeline and how to run it locally C. The C2DAO branch workflow and promotion requirements D. CBAC configuration, tenant isolation architecture, and the absolute prohibition on CBAC bypass — because a misunderstanding of tenant isolation is the failure mode with the highest potential for irreversible data exposure harm

15. A platform architecture review reveals that two separate MSS programs are querying the same large Vehicle Object set independently, each running full-table scans 12 times per day. The correct SL 5L architectural recommendation is:

A. Implement a shared caching layer (or a pre-computed Foundry dataset view) that both programs consume — eliminating redundant full-table scans and reducing platform load B. Migrate both programs to use gRPC instead of the OSDK for better query performance C. Add indexes to all queried properties on the Vehicle Object Type D. Reduce each program's query frequency to 6 times per day as a rate-limiting measure

---

SECTION 2 — SHORT ANSWER

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

SA-1. You are leading a security assessment of a new OSDK-based TypeScript integration that allows division S3 staff to query vehicle readiness data and submit maintenance status updates via Actions. Describe the full security assessment methodology: what you check, what tools you use, and the minimum findings that would block release.

SA-2. The USAREUR-AF MSS platform team is asked to design a high-throughput ingestion architecture for a new theater logistics system that will push 30,000 supply requisition records per minute via an external feed. Describe the end-to-end architecture: ingestion protocol, buffering/batching strategy, transaction pattern, input validation, and how you would prevent malformed external data from corrupting the Ontology.

SA-3. A senior engineer submits a pull request that modifies the CBAC policy for the 1st Armored Division tenant — expanding read access to a new set of Object Types. Describe the complete review and approval process required before this change can merge to production, citing all required SL 5L governance steps.

SA-4. You are asked to design the DevSecOps pipeline for the USAREUR-AF MSS Ontology CI system. Describe each stage of the pipeline, what it validates or tests, what tools you would use, and what the blocking criteria are for each stage (what failures block the merge).

SA-5. Three months after a new Platform SDK-based pipeline went into production, a data steward reports that the 1st Cavalry Division's Vehicle readiness data briefly appeared in a query run by a 3rd Armored Brigade Combat Team analyst who should not have access. Conduct a root-cause analysis: what failure controls should have caught this, what is the most likely architectural cause, and describe the remediation and reporting steps required under SL 5L.

SCORING SUMMARY

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

Passing: 42/60 (70%) — Post-test only. Pre-test is diagnostic.

ANSWER KEY — INSTRUCTOR USE ONLY

Do not distribute to students.

Multiple Choice: 1. C — CBAC bypass is NEVER acceptable in any environment; set up properly scoped test data instead. 2. B — Report immediately as security incident, rotate credentials, audit access log, investigate root cause. 3. B — Ontology CI validates that downstream consumers (OSDK queries, validators, FOOs, pipelines) will not break. 4. D — C2DAO sign-off required for both Ontology branch merges AND CBAC policy changes. 5. A — Buffer into batches, write in atomic transaction, advance checkpoint only after successful commit. 6. C — Profile at each stage (client pagination, query execution, network) before applying any optimization. 7. C — Cache-aside: fetch from Ontology on miss, populate cache, serve from cache on hit. 8. D — Security review of external data feed required (provenance, schema, input validation) before activation. 9. C — Production profiling requires C2DAO data steward coordination to prevent degrading operational users. 10. B — Critical injection vulnerability — unsanitized user input in query logic can expose unauthorized records. 11. B — Peer reviewer must be SL 5L qualified to review platform-level CBAC, Ontology CI, and SDK integrations. 12. A — AR 25-2 (Army Cybersecurity) governs protection from unauthorized data access; cross-tenant bleed = violation. 13. A — Branch merges to production require C2DAO human sign-off — not automated regardless of check results. 14. D — CBAC and tenant isolation first — highest potential harm from misunderstanding; CBAC bypass can expose operational data irreversibly. 15. A — Shared caching layer or pre-computed view eliminates redundant full-table scans across both programs.

Short Answer Guidance:

SA-1. Full credit: methodology — (1) SAST: run static analysis on the TypeScript codebase for injection vulnerabilities (unsanitized input in queries/Actions), hardcoded credentials, overly broad permissions; (2) DAST: test the running integration with adversarial inputs — malformed unit designations, over-length strings, injection attempts in Action parameter values; (3) CBAC verification: confirm that queries only return records within the authenticated user's authorized scope; (4) input sanitization audit: review all points where user-supplied strings are used in queries or Action parameters; (5) credential handling review: no SDK credentials in code, all injected via environment; blocking findings: any injection vulnerability in query/Action paths, hardcoded credentials, CBAC misconfiguration allowing out-of-scope record access, unhandled API errors that expose internal details. All five methodology steps and at least three blocking criteria required.

SA-2. Full credit: ingestion protocol — gRPC (high-throughput, low-latency; 30K records/min requires binary serialization); buffering/batching — consumer buffers incoming events into batches of configurable size (e.g., 1,000 records or 5 seconds, whichever comes first); transaction pattern — each batch written as an atomic Foundry write transaction; Kinesis/stream checkpoint advanced only after successful commit — if transaction fails, batch is retried without data loss; input validation — all incoming records validated against a strict schema (required fields, type checks, value range checks) BEFORE being included in the write transaction; records that fail validation are routed to an error queue for investigation rather than dropped silently or written; preventing Ontology corruption — schema validation layer between stream consumer and Foundry write; no external field passes directly into the Ontology without explicit type-checked mapping. All five components required for full credit.

SA-3. Full credit: required steps in order — (1) CBAC policy changes require peer review by a second SL 5L qualified engineer (not just any reviewer); (2) the second SL 5L engineer reviews the policy change for correctness and potential cross-tenant exposure risk; (3) submit a C2DAO branch for the CBAC change with a complete description of what access is being expanded, why, and what data is now accessible; (4) C2DAO data steward reviews the CBAC change independently — this is a separate human sign-off from the peer review; (5) data steward approves the branch merge; (6) after merge, verify the CBAC change in staging by confirming that (a) the intended access is now permitted and (b) access that should remain restricted is still denied; (7) document the change in the access control configuration log. All seven steps required; peer review by SL 5L qualified engineer and C2DAO sign-off are both mandatory.

SA-4. Full credit: pipeline stages — (1) Linting and format check: ensures code conforms to platform style standards; blocking: any lint error; (2) SAST scan: static analysis for security vulnerabilities (injection, hardcoded secrets, broad permissions); blocking: any HIGH or CRITICAL finding; (3) Unit tests: validates individual functions and modules; blocking: any test failure; (4) Ontology schema compatibility check: validates that renamed or removed properties are not referenced by any downstream consumer (OSDK query, Action validator, FOO, Pipeline Builder); blocking: any breaking schema change without migration path; (5) Integration tests: validates end-to-end behavior in a staging environment with scoped test data; blocking: any test failure or CBAC misconfiguration detected; (6) Security scan for dependency vulnerabilities: checks package dependencies for known CVEs; blocking: any CRITICAL CVE in a direct

dependency; (7) C2DAO promotion review (human step): data steward reviews change description and approves merge. All stages with tools and blocking criteria required; human C2DAO step must be identified as the final gate.

SA-5. Full credit: root cause analysis — most likely architectural cause: CBAC policy misconfiguration — either the 1st Cavalry Division data does not have a `owning_unit` attribute set correctly, or the 3rd ABCT user's CBAC policy was configured with an overly broad grant; failed controls: (1) CBAC configuration peer review by second SL 5L engineer should have caught the overly broad policy; (2) Ontology CI should include a CBAC policy test that validates out-of-scope access is denied; (3) staging environment access testing should have verified cross-tenant isolation before production deployment; remediation steps: (1) immediately audit the CBAC configuration and restrict access to correct scope; (2) audit the access log for the full exposure window to determine what data was accessed and by whom; (3) report the incident through the security incident reporting process (AR 25-2 requires reporting of unauthorized access); (4) notify the 1st Cavalry Division data owner of the exposure; (5) fix the CBAC policy with C2DAO sign-off; (6) add a regression test to Ontology CI that validates this specific cross-tenant boundary. Must include exposure audit, security incident report requirement, and regression test addition.

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