PUBLICATION

EXAM-TM50O-POST

POST-TEST — SL 5O: ADVANCED PLATFORM 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 5O: ADVANCED PLATFORM ENGINEER

MAVEN SMART SYSTEM (MSS) — USAREUR-AF

Field	Detail
Course	SL 5O: Advanced Platform Engineer
Form	Post-Test
Level	SL 5O (Advanced Specialist)
Audience	Experienced platform engineers completing SL 5O training
Time Allowed	45 minutes
Passing Score	70% (39/56)

INSTRUCTIONS

This assessment evaluates knowledge and skills gained during SL 5O training. Answer all questions. A score of 70% or higher is required for course credit.

SECTION 1 — MULTIPLE CHOICE

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

1. In a wave-based fleet upgrade strategy, the correct wave order is:

A. Production clusters first, then dev/test, then edge B. All clusters simultaneously to minimize version skew C. Edge clusters first because they have the fewest users D. Dev/test → canary production → regional hubs → edge clusters → high-sensitivity clusters

2. When the error budget is exhausted (0% remaining), the correct response per SRE practice is:

A. Change freeze for non-critical changes; dedicate engineering time to reliability improvements; require incident review before resuming deployments B. Continue deploying features but increase monitoring C. Lower the SLO target so the budget is no longer exhausted D. Stop all monitoring to prevent further budget consumption

3. "Continuous compliance" differs from traditional compliance because:

A. It requires fewer security controls B. It monitors control effectiveness continuously and generates evidence from live system data — not periodic manual assessments C. It eliminates the need for an ATO D. It only applies to cloud environments

4. A "golden path" in platform engineering is:

A. The only allowed way to deploy applications B. A physical network path with guaranteed bandwidth C. An opinionated, pre-configured, tested, and documented approach for common tasks — not mandatory, but supported and recommended D. A security classification level between Secret and Top Secret

5. Federated observability across a fleet means:

A. Each cluster has its own isolated monitoring with no cross-cluster visibility B. Metrics, logs, and traces from all clusters are aggregated or queryable from a central point, enabling cross-cluster correlation C. Only the hub cluster has monitoring D. Edge clusters send all raw data to the hub in real time

6. In the incident management framework, "blameless post-incident review" means:

A. No one is held accountable for incidents B. The review is conducted anonymously C. Incidents are not documented D. The review focuses on contributing factors, system design, and process improvements — not on individual blame — because blame discourages honesty and prevents learning

7. DORA metrics include all of the following EXCEPT:

A. Deployment frequency B. Lead time for changes C. Lines of code per developer per sprint D. Time to restore service

8. When alerting at fleet scale, the principle "alert on symptoms, not causes" means:

A. Alert on user-facing impact (SLO burn rate, error rate, latency) rather than internal system metrics (CPU, memory, disk) — because symptoms tell you something is broken for users, causes may not B. Only alert when users complain C. Disable all alerts and rely on dashboards D. Only alert on hardware failures

SECTION 2 — SHORT ANSWER

Answer in 2–3 sentences. (5 points each)

9. Explain the difference between "hub" and "edge" clusters in the MSS fleet topology. What assumptions differ about connectivity, capacity, and autonomy?

10. Describe how policy-as-code STIG automation works. What happens when a check fails? What happens when a legitimate exception exists?

11. What is the risk of upgrading all clusters in a fleet simultaneously instead of using a wave strategy?

12. Explain why "developer experience" is a platform engineering responsibility. How do DORA metrics connect platform quality to developer productivity?

SECTION 3 — SCENARIO

Answer in 5–8 sentences. (10 points each)

13. You are responsible for the MSS fleet upgrade from Kubernetes 1.29 to 1.30. Your fleet has 3 dev/test clusters, 1 canary production cluster, 4 regional hub clusters, and 12 edge clusters. Describe your upgrade plan: wave structure, validation gates between waves, rollback criteria, and how you would handle an edge cluster that is disconnected during its upgrade window.

14. The ISSM reports that the MSS ATO annual review is in 60 days and asks for compliance evidence across all clusters. Under the traditional process, this would take two engineers 4 weeks of manual evidence gathering. Describe how SL 5O continuous compliance automation addresses this: what evidence is already being generated, how is it stored, and how would you produce the ATO package.

SCORING SUMMARY

Section	Questions	Points Each	Total Points
Multiple Choice	8	2	16
Short Answer	4	5	20
Scenario	2	10	20
Total	—	—	**56**

Passing: 39/56 (70%) — Post-test only. Pre-test is diagnostic.

ANSWER KEY — INSTRUCTOR USE ONLY

Do not distribute to students.

Multiple Choice: 1. D — Wave order: dev/test → canary production → regional hubs → edge clusters → high-sensitivity clusters (lowest risk first). 2. A — Exhausted error budget triggers a change freeze for non-critical changes, dedicated reliability engineering time, and incident review before resuming. 3. B — Continuous compliance monitors control effectiveness continuously from live system data, not periodic manual assessments. 4. C — A golden path is an opinionated, pre-configured, tested, and documented approach — not mandatory, but supported and recommended. 5. B — Federated observability aggregates metrics, logs, and traces from all clusters, enabling cross-cluster correlation from a central point. 6. D — Blameless review focuses on contributing factors, system design, and process improvements — blame discourages honesty and prevents learning. 7. C — Lines of code per developer per sprint is NOT a DORA metric. The four DORA metrics are: deployment frequency, lead time for changes, time to restore service, and change failure rate. 8. A — Alert on user-facing impact (SLO burn rate, error rate, latency) rather than internal system metrics, because symptoms confirm user impact.

Short Answer Guidance:

SA-9. Full credit: hub clusters are centrally located with reliable connectivity, higher capacity (compute/storage), and serve as aggregation and management points — they run centralized services (observability, GitOps controllers, artifact registries). Edge clusters are deployed forward with constrained connectivity (DDIL), limited capacity, and must operate autonomously during disconnection — they cache images, queue telemetry, and reconcile with the hub when connectivity returns. Key assumption differences: hub assumes stable connectivity and sufficient resources; edge assumes intermittent connectivity, limited resources, and the need for autonomous operation. Partial credit (3 pts) for correct distinction without connectivity/autonomy assumptions.

SA-10. Full credit: policy-as-code STIG automation encodes each STIG control as an executable check (e.g., Open Policy Agent, Kyverno, or InSpec rules) that runs continuously against cluster configuration. When a check fails: the system generates a compliance finding, alerts the security team, and can optionally block the non-compliant resource from deploying (admission control). When a legitimate exception exists: the exception is documented as code (exception policy with justification, approver, and expiration date), version-controlled alongside the policy, and auditable — the check still runs but the finding is marked as "accepted risk" rather than suppressed. Partial credit (3 pts) for correct automation description without exception handling.

SA-11. Full credit: simultaneous upgrade risks: (1) a bug in the new version affects ALL clusters at once — no healthy clusters to fall back on; (2) if rollback is needed, all clusters must roll back simultaneously, extending the outage; (3) version-specific issues may only manifest under certain workloads — a wave strategy catches these on low-risk clusters before they reach production; (4) operational impact: if every

cluster is upgrading at once, there may be no operational environment available during the upgrade window. A wave strategy limits blast radius — if Wave 1 (dev/test) fails, production is unaffected. Partial credit (3 pts) for identifying one risk without connecting to wave strategy mitigation.

SA-12. Full credit: developer experience is a platform engineering responsibility because the platform's "users" are developers — if the platform is hard to use, developers will work around it (shadow IT), creating security and compliance gaps. DORA metrics connect: (1) lead time for changes — a good platform (golden paths, self-service, automated CI/CD) reduces lead time; a bad platform increases it with manual steps and waiting; (2) deployment frequency — self-service deployment pipelines enable more frequent, smaller deployments; (3) change failure rate — golden paths with built-in testing and security reduce failures; (4) time to restore — good observability and rollback mechanisms (platform responsibilities) reduce MTTR. Partial credit (3 pts) for connecting DX to platform without DORA metric specifics.

Scenario Guidance:

S-13. Full credit (10 pts): upgrade plan: Wave 1 — 3 dev/test clusters (validate basic functionality, run integration tests); validation gate: all smoke tests pass, no elevated error rates for 24 hours. Wave 2 — 1 canary production cluster (real production traffic, real workloads); validation gate: SLO metrics stable for 48 hours, no increase in error rate or latency. Wave 3 — 4 regional hub clusters (one at a time or in pairs); validation gate: SLO stable for 24 hours per hub, federated observability confirms cross-cluster health. Wave 4 — 12 edge clusters (batched by region/priority); validation gate: health check reports from each edge cluster. Rollback criteria: any wave that shows SLO degradation beyond a defined threshold (e.g., error rate >2x baseline) triggers automatic rollback of that wave and halt of subsequent waves. Disconnected edge cluster: skip it in the current wave, document the deferral, set a reconciliation flag — when the cluster reconnects, it enters a deferred upgrade queue and is upgraded as a solo wave with its own validation gate. Must cover wave structure, gates, rollback, and disconnected handling. Partial credit (5 pts) for wave structure without gates or disconnected handling.

S-14. Full credit (10 pts): continuous compliance automation means: (1) policy-as-code checks run continuously against all clusters — every STIG control is an automated check that generates evidence (pass/fail, timestamp, resource, cluster ID) stored in a compliance database; (2) admission controllers prevent non-compliant resources from deploying — so evidence of compliance is generated at deploy time; (3) configuration drift detection flags any manual changes that deviate from declared state; (4) to produce the ATO package: query the compliance database for the last 12 months of evidence across all clusters, generate the report automatically — this should take hours, not weeks; (5) gaps: any controls that cannot be fully automated have a documented manual check process with scheduled evidence collection. The key advantage: the evidence already exists because it is generated continuously — the ATO review is an export, not a collection exercise. Must describe what evidence is already being generated, how it is stored, and how the package is produced. Partial credit (5 pts) for describing automation without explaining the ATO package production process.