


DRAFT — UNOFFICIAL — NOT FOR OPERATIONAL USE

PUBLICATION

EXAM-TM500-PRE



PRE-TEST — SL 50: ADVANCED PLATFORM 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

PRE-TEST — SL 50: ADVANCED PLATFORM ENGINEER

MAVEN SMART SYSTEM (MSS) — USAREUR-AF

Field	Detail
Course	SL 50: Advanced Platform Engineer
Form	Pre-Test
Level	SL 50 (Advanced Specialist)
Audience	Experienced platform engineers; prerequisite: SL 40 Go
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. "Cluster API" in Kubernetes is used to:

A. Provide an API for applications running on the cluster
B. Manage the lifecycle of Kubernetes clusters themselves — provisioning, upgrading, and decommissioning clusters declaratively
C. Create API gateways for external traffic
D. Monitor API response times

2. An "SLO" (Service Level Objective) defines:

A. The maximum amount of money to spend on infrastructure B. The minimum number of engineers required to operate a service C. A target level of reliability for a service, measured by specific indicators (SLIs), with a defined error budget D. The service's feature roadmap for the next quarter

3. An "error budget" is:

A. The acceptable amount of unreliability — the gap between 100% and the SLO target — used as a decision framework for balancing velocity and reliability B. The amount of money allocated for fixing bugs C. The number of error messages an application is allowed to display D. The maximum number of incidents per month

4. "Policy-as-code" for STIG compliance means:

A. Writing STIG documentation in a programming language B. Using code comments to document which STIGs apply C. Hiring developers to manually review STIG checklists D. Encoding security requirements as machine-executable policies that automatically check system configuration and report compliance

5. The DORA "lead time for changes" metric measures:

A. How long it takes to hire a new developer B. The time from a code commit to that code running in production C. How long it takes to write a design document D. The time between product releases

SECTION 2 — SHORT ANSWER

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

6. What is a "golden path" in platform engineering? Why is it valuable even though developers are not required to follow it?

7. Explain the difference between managing a single Kubernetes cluster (SL 40) and managing a fleet of clusters (SL 50). What new challenges emerge at fleet scale?

8. What is a "blameless post-incident review"? Why is "blameless" important?

SECTION 3 — SCENARIO

Answer in 3–5 sentences. (10 points)

9. Your SLO for application availability is 99.9% (43 minutes of downtime per month). It is day 15 of the month and you have already consumed 35 minutes of error budget. A product team wants to deploy a large new feature. What do you recommend, and why?

SCORING SUMMARY

Section	Questions	Points Each	Total Points
Multiple Choice	5	2	10
Short Answer	3	5	15
Scenario	1	10	10
Total	—	—	35

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

ANSWER KEY — INSTRUCTOR USE ONLY

Do not distribute to students.

Multiple Choice: 1. B — Cluster API manages the lifecycle of Kubernetes clusters themselves — provisioning, upgrading, and decommissioning declaratively. 2. C — An SLO is a target reliability level measured by SLIs (Service Level Indicators) with a defined error budget. 3. A — The error budget is the acceptable unreliability gap between 100% and the SLO target, used to balance velocity and reliability decisions. 4. D — Policy-as-code encodes security requirements as machine-executable policies that automatically check configuration and report compliance. 5. B — Lead time for changes measures the time from code commit to that code running in production.

Short Answer Guidance:

SA-6. Full credit: a golden path is an opinionated, pre-configured, tested, and documented approach for common tasks (e.g., "deploy a new Python service" template with CI/CD, monitoring, RBAC, and network policies pre-configured). It is valuable because: (1) it encodes best practices — developers get security, observability, and compliance by default without being experts in each; (2) it reduces onboarding time — new teams can ship a working, compliant service in hours instead of weeks; (3) it is not mandatory — developers can deviate, but the golden path is the supported, recommended, and easiest option. Partial credit (3 pts) for correct definition without explaining why it is valuable despite being optional.

SA-7. Full credit: single cluster (SL 40) = managing one Kubernetes environment — deployments, networking, storage, RBAC, and security within that cluster. Fleet (SL 50) = managing many clusters across environments (dev, prod, edge, air-gapped) — new challenges: configuration consistency across clusters, version skew management, fleet-wide upgrades (wave strategy), centralized observability across clusters, policy enforcement at scale, and handling clusters with different connectivity (hub vs. edge). Partial credit (3 pts) for correct distinction without naming specific fleet-scale challenges.

SA-8. Full credit: a blameless post-incident review analyzes an incident by focusing on contributing factors, system design, and process improvements — not on identifying individuals to blame. "Blameless" is important because: (1) blame discourages honesty — engineers will hide mistakes or avoid reporting near-misses if they fear punishment; (2) the root cause of incidents is almost always a system/process failure, not an individual failure — blaming individuals misses the systemic fix; (3) blameless culture encourages proactive reporting, which leads to faster detection and prevention of future incidents. Partial credit (3 pts) for correct definition without explaining why blamelessness matters.

Scenario Guidance:

S-9. Full credit (10 pts): recommendation: do not deploy the large feature now — with 35 of 43 minutes consumed by day 15, only 8 minutes of error budget remain for the rest of the month; a large feature deployment carries elevated risk of incidents. Instead: (1) delay the deployment to the start of the next month when error budget resets; (2) if the feature is operationally urgent, require additional risk mitigation — canary deployment to a small subset, feature flags for instant rollback, enhanced monitoring during rollout; (3) use this as a trigger for a reliability review — why was 81% of the budget consumed in 50% of the month? Are there systemic reliability issues to address? Must cite the specific error budget numbers, recommend a course of action, and connect the recommendation to SRE error budget principles. Partial credit (5 pts) for correct recommendation without error budget analysis.

USAREUR-AF Operational Data Team TM-500 Pre-Test | Version 1.0 | March 2026