

DRAFT — UNOFFICIAL — NOT FOR OPERATIONAL USE

PRACTICAL EXERCISE

EX-500



EX_500 — Advanced Platform Engineer

Practical Exercise — SL 50 Proficiency

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

EX_500 — ADVANCED PLATFORM ENGINEER

PRACTICAL EXERCISE — SL 50 PROFICIENCY

Field	Value
Version	1.0 — March 2026
Prerequisite	SL 40 REQUIRED (Go evaluation on file); SL 50 — Advanced Platform Engineer Technical Manual
Duration	3–4 hours
Environment	Multi-cluster training environment with fleet management tooling — see EX_400 ENVIRONMENT_SETUP.md (extended environment)

COMPANION RESOURCES

Resource	Reference
Technical Manual	SL 50 — Advanced Platform Engineer
Syllabus	SYLLABUS_TM500
Pre-Exercise Exam	EXAM_TM500_PRE
Post-Exercise Exam	EXAM_TM500_POST

SCENARIO

MSS has expanded beyond a single cluster to a fleet spanning hub and edge locations across the USAREUR-AF AOR. The OPDATA team needs a fleet management strategy, SLO framework, compliance automation, and a golden path for application onboarding. The trainee will design and implement fleet-level platform capabilities.

TASK LIST

Task 1 — Fleet Topology and Upgrade Strategy (45 min)

- Design a fleet topology: hub clusters, edge clusters, management plane, connectivity assumptions
- Define a wave-based upgrade strategy with at least 3 waves (dev → canary prod → regional hubs → edge)
- Specify validation gates between waves (soak time, health checks, SLO verification)
- Document rollback procedure for a failed upgrade at each wave

Standard	Criteria
Go	Topology includes hub and edge with connectivity assumptions; 3+ upgrade waves with validation gates; rollback documented per wave
No-Go	No fleet topology; fewer than 3 waves; no validation between waves; no rollback procedure

Task 2 — SLO Framework (30 min)

- Define SLOs and SLIs for at least 3 MSS platform services (e.g., application availability, CI/CD pipeline, data freshness)
- Calculate error budgets for each SLO
- Define error budget policy: what happens at 50%, 25%, and exhausted thresholds

Standard	Criteria
Go	3+ SLOs with measurable SLIs; error budgets calculated; budget policy has actionable thresholds
No-Go	Fewer than 3 SLOs; SLIs not measurable; no error budget policy

Task 3 — Compliance Automation (45 min)

- Implement a policy-as-code check for at least 3 STIG-equivalent controls (e.g., container runs non-root, resource limits set, network policy exists)
- Build a compliance dashboard showing pass/fail/exception status per control
- Generate an evidence export suitable for RMF documentation

Standard	Criteria
Go	3+ automated policy checks; dashboard shows real status; evidence export is auditable

Standard	Criteria
No-Go	Manual checks only; no dashboard; evidence not exportable

Task 4 — Post-Incident Review (30 min)

- Given a provided incident scenario (fleet-wide deployment causes service degradation across 3+ clusters), produce a blameless post-incident review document
- Document the incident timeline using all six phases: detect, triage, mitigate, resolve, review, improve
- Identify at least 3 contributing factors (not root cause alone) and explain why each contributed to the incident
- Define at least 3 action items with owners, priorities, and due dates that address the contributing factors
- Explain how the post-incident review feeds back into the SLO framework (error budget impact, policy threshold changes)

Standard	Criteria
Go	Blameless tone throughout; all 6 phases documented with timeline; 3+ contributing factors identified; 3+ action items with owners and deadlines; SLO/error budget connection explicit
No-Go	Blame assigned to individuals; phases missing from timeline; fewer than 3 contributing factors; action items lack owners or deadlines; no connection to SLO framework

Task 5 — Fleet Observability (30 min)

- Design a federated observability architecture covering metrics, logs, and traces across the fleet topology from Task 1
- Configure cross-cluster metric federation with a hierarchical collection strategy (edge → hub → fleet aggregator)
- Implement SLO-based alerting for at least 2 of the SLOs defined in Task 2 using burn-rate alerts instead of static thresholds
- Define an alert severity model (P1/P2/P3) with routing rules and escalation procedures per tier
- Demonstrate cross-cluster correlation: given a symptom on one cluster, show how the observability stack traces impact across the fleet

Standard	Criteria
Go	All three pillars (metrics, logs, traces) addressed; federation strategy matches fleet topology; 2+ SLO burn-rate alerts configured; severity model with routing defined; cross-cluster correlation demonstrated
No-Go	Any observability pillar missing; no federation strategy; alerts use static thresholds only; no severity model; no cross-cluster correlation

Task 6 — Golden Path and Developer Experience (30 min)

- Design a golden path for MSS application onboarding: define the template, pre-configured pipeline, monitoring, and security scanning
- Document the golden path steps from `create-mss-app` to first production deployment
- Define DORA metrics targets for the platform and explain how the golden path supports each metric

Standard	Criteria
Go	Golden path is end-to-end (template → deploy); DORA metrics defined with targets; golden path demonstrates how it improves each metric
No-Go	Golden path is incomplete; no metrics; path does not connect to measurable outcomes