


DRAFT — UNOFFICIAL — NOT FOR OPERATIONAL USE

PUBLICATION

# EXAM-TM40G-POST



---

## POST-TEST — SL 4G: ORSA

---

*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

## POST-TEST — SL 4G: ORSA

### MAVEN SMART SYSTEM (MSS) — USAREUR-AF

Field	Detail
Course	SL 4G: ORSA
Form	Post-Test
Level	SL 4G (Specialist)
Audience	FA49 / ORSA analysts; prerequisite: SL 1+20+30 + graduate-level quantitative methods + Python or R proficiency
Time Allowed	45 minutes
Passing Score	70% (46/66)

### 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. When writing output from a Code Workspace analysis back to a Foundry curated dataset, the correct pattern is:**

- A. Use a write transaction to commit results directly to the authorized output dataset via the Foundry SDK
- B. Export results to CSV and email to the data steward for manual upload
- C. Write results to a temporary file in the workspace and schedule a pipeline to pick them up
- D. Push results to a shared Git repository

where a pipeline reads them on the next run

**2. A readiness regression model you built shows  $R^2 = 0.91$  on training data but  $R^2 = 0.52$  on a held-out test set. The MOST likely explanation and correct response is:**

A. The test set is corrupted — re-run the pipeline with a different random split B. The  $R^2$  gap is expected for military data; proceed to production C. The model is acceptable —  $R^2 = 0.91$  is sufficient and test set performance will improve with more data D. The model is overfitting the training data; apply regularization or reduce model complexity and re-evaluate

**3. Per USAREUR-AF ORSA product standards, every point estimate presented to a commander MUST be accompanied by:**

A. A footnote explaining the data source B. A comparison to the previous year's value C. A confidence or prediction interval with the confidence level explicitly stated D. A p-value demonstrating statistical significance

**4. You are selecting an ARIMA model for a weekly maintenance backlog time series. You observe that the ACF plot shows gradual decay and the PACF has a sharp cutoff after lag 2. This pattern suggests:**

A. An MA(2) model — moving average order 2 B. An AR(2) model — autoregressive order 2 C. An ARIMA(0,1,2) model requiring differencing D. A seasonal ARIMA model

**5. The minimum number of Monte Carlo simulation trials required under USAREUR-AF ORSA standards to estimate COA risk at defined probability thresholds is:**

A. 100 trials B. 1,000 trials C. 500 trials D. 10,000 trials

**6. You are conducting a Monte Carlo simulation of two COAs. COA 1 shows a 15% probability of mission failure; COA 2 shows a 22% probability but higher expected gain. How should you present this finding to the commander?**

A. Recommend COA 1 only — lower risk is always preferred in Army operations B. Present both COAs with their full probability distributions of outcomes, explicitly state assumptions, and let the commander decide based on operational context C. Average the two COAs and present the average as the recommended course of action D. Drop COA 2 from the analysis — a 22% failure rate exceeds acceptable risk

**7. A linear programming model for vehicle allocation has an objective function that maximizes total readiness. After solving, the shadow price for the maintenance man-hours constraint is 0.004 readiness units per hour. This means:**

A. The model is infeasible — the shadow price should be zero B. Each additional man-hour beyond the current constraint limit would increase readiness by 0.004 units at the margin C. The current solution uses 0.4% of the available man-hours D. The constraint is not binding and can be removed from the formulation

**8. Per USAREUR-AF ORSA product standards, a simulation's random seed must be:**

A. Set to the current date to ensure temporal reproducibility B. Changed for each run to prevent deterministic outputs C. Documented in the product so any analyst can reproduce the exact results D. Generated by the Foundry platform to ensure platform-level reproducibility

**9. Your S3 asks for a demand forecast. You provide a forecast with 90% confidence bounds. The S3 says "90% is not certain enough — can you make it 99%?" The correct response is:**

A. Re-run the model with a higher confidence level parameter set to 99% B. Use both a 90% and 99% CI in the final product and let the commander choose C. Collect more data until the 90% CI is narrow enough to effectively function as a 99% CI D. Explain that widening to 99% CI produces a wider interval with less operational precision; confirm the appropriate level with the S3 based on the decision being made

**10. In a commander's brief using MSS Contour or Quiver visualizations, uncertainty bounds should be displayed as:**

A. A footnote at the bottom of the chart B. A separate appendix slide titled "Model Uncertainty" C. Error bars, shaded confidence bands, or explicitly labeled min/max bounds directly on the chart D. A verbal caveat during the brief, without charting the bounds

**11. You build a logistic regression model to predict whether a vehicle will fail inspection. The model achieves 88% accuracy on the test set. However, you notice that 85% of all vehicles pass inspection. Which additional metric is MOST important to evaluate?**

A. R-squared B. Mean absolute error C. Precision, recall, and the F1 score for the failure class D. AIC/BIC model selection criteria

**12. When connecting a Python Code Workspace to Foundry datasets, the correct pattern for reading an input dataset is:**

A. Use the Foundry dataset connector to read the dataset directly within the workspace via the approved SDK B. Download the dataset as a CSV from the MSS UI and read it with `pd.read_csv()` C. Clone the dataset's Git repository and import it as a Python module D. Submit a data request to the data steward who exports the data and emails it to you

**13. An ORSA team is asked to support a COA comparison. COA A has a higher expected readiness gain; COA B has lower variance in outcomes. Per SL 4G, the analyst's role in the brief is to:**

A. Present both COAs with their quantified risk profiles and let the commander make the decision based on operational risk tolerance B. Recommend the COA with higher expected gain — that is always the analytically superior choice C. Recommend COA B — the Army doctrine preference is risk mitigation D. Present only the metric on which the two COAs differ most

**14. You are presenting a linear programming result to a G4. The G4 asks: "What if we could get 50 more maintenance hours per week?" The analytical tool used to answer this question is:**

A. A new Monte Carlo simulation run B. A bootstrapped confidence interval for the optimal solution C. A revised regression model trained on additional data D. The LP sensitivity analysis / shadow price report

**15. Per USAREUR-AF ORSA product standards, the documentation package for an analytical product must include:**

A. Methodology description, data sources, assumptions, limitations, validation statistics, and reproducibility information (including seed) B. Model code only — the commander's brief serves as the written documentation C. An executive summary and a copy of the raw input data D. A list of ORSA analysts who reviewed the product

---

## SECTION 2 — SHORT ANSWER

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

**SA-1. You have built a time series forecast of equipment failure rates for a division maintenance tracker. Describe how you would validate the model, what statistics you would report, and how you would communicate the model's limitations to the G4 in an operational brief.**

**SA-2. A Monte Carlo simulation comparing two COAs shows: COA A — 10th percentile outcome: -12 units readiness, 50th: +8, 90th: +22. COA B — 10th: +2, 50th: +9, 90th: +14. The S3 says "COA A has the higher upside, recommend it." Write the ORSA response you would provide to the S3, applying SL 4G standards.**

**SA-3. Describe the complete write transaction workflow for committing ORSA model outputs from a Code Workspace back to a Foundry curated dataset. Include the purpose of each step.**

**SA-4.** You are asked to optimize maintenance man-hour allocation across eight units to maximize total fleet readiness, subject to: a total man-hour cap of 800 hours/week, a minimum of 50 hours per unit, and a constraint that no unit receives more than 150 hours. Describe the LP formulation: decision variables, objective function, and constraints.

**SA-5.** Explain what is wrong with the following commander's brief statement, and rewrite it to meet USAREUR-AF ORSA product standards: "The model predicts the unit will have 94% readiness next month."

**SA-6.** Describe how an ORSA analyst's quantitative work directly supports two WFF functions. For each, identify the WFF track (SL 4A through SL 4F) and give a concrete example of an ORSA product that supports decision-making in that function.

## SCORING SUMMARY

Section	Questions	Points Each	Total Points
Multiple Choice	15	2	30
Short Answer	6	6	36
<b>Total</b>	—	—	<b>66</b>

Passing: 46/66 (70%) — Post-test only. Pre-test is diagnostic.

## ANSWER KEY — INSTRUCTOR USE ONLY

*Do not distribute to students.*

**Multiple Choice:** 1. A — Write transaction via Foundry SDK is the correct Code Workspace output pattern. 2. D — Large gap between train and test  $R^2$  is the classic overfitting signature; apply regularization. 3. C — Every point estimate must include a CI with confidence level stated — USAREUR-AF standard. 4. B — Gradual ACF decay + PACF cutoff at lag 2 = AR(2) signature. 5. B — USAREUR-AF ORSA standard: minimum 1,000 trials for Monte Carlo. 6. B — Present full distributions, state assumptions, defer COA decision to commander. 7. B — Shadow price = marginal value of relaxing the constraint by one unit. 8. C — Seed must be documented for reproducibility. 9. D — Wider CI = wider interval, less operational precision; confirm appropriate level with decision-maker. 10. C — Uncertainty bounds displayed directly on charts (error bars, shaded bands, labeled ranges). 11. C — Class imbalance makes accuracy misleading; precision/recall/F1 for the minority (failure) class is essential. 12. A — Foundry dataset connector via approved SDK is the correct Code Workspace read pattern. 13. A — ORSA presents quantified risk profiles; COA decision rests with the commander. 14. D — LP sensitivity analysis / shadow price answers "what if constraint changes" questions. 15. A — Full documentation package required: methodology, sources, assumptions, limitations, validation, reproducibility.

**Short Answer Guidance:**

SA-1. Full credit: validation approach — hold out last N periods, calculate RMSE/MAE/MAPE on holdout; report RMSE, MAE, and forecast interval coverage; limitations for G4 brief: model is trained on historical data and may not capture new failure modes, confidence intervals widen at longer horizons, model should be re-validated when operational tempo changes. Must include validation statistics AND a specific limitation for full credit.

SA-2. Full credit: ORSA response should note that COA A has significant downside risk (10th percentile = -12 readiness units) while COA B has a positive floor (10th percentile = +2); the choice between upside potential and downside protection is a commander's decision based on operational risk tolerance, not an analytical recommendation; ORSA should present both distributions and explicitly state: "ORSA presents these results for the commander's decision. COA A offers higher upside but exposes the force to significant downside at the 10th percentile. COA B offers a more stable outcome range." Full credit requires citing downside risk AND deferring decision to commander.

SA-3. Full credit: (1) complete all analysis and generate output DataFrame in the workspace; (2) open a write transaction on the target output dataset via the Foundry SDK; (3) write the DataFrame within the transaction context; (4) commit (close) the transaction — this atomically publishes the results; purpose: the transaction ensures the write is atomic (either all records write or none), preventing partial writes that could corrupt downstream pipelines. Full credit requires transaction pattern AND purpose explanation.

SA-4. Full credit: decision variables:  $x_i$  = man-hours allocated to unit  $i$  ( $i = 1..8$ ); objective: maximize  $\sum f(x_i)$  where  $f$  is a readiness function of man-hours (or proxy); constraints:  $\sum x_i \leq 800$  (total cap),  $x_i \geq 50$  for all  $i$  (minimum per unit),  $x_i \leq 150$  for all  $i$  (maximum per unit),  $x_i \geq 0$ . Full credit requires all three constraint types plus decision variables and objective. Partial credit (3 pts) for correct decision variables and objective without all constraints.

SA-5. Full credit: the statement has two problems — (1) it presents only a point estimate without a confidence interval; (2) "the model predicts" implies certainty; corrected statement: "Analytical estimate: 94% readiness (90% CI: 88%–98%). This estimate is based on [model type], trained on [time period], and assumes [key assumption]. Limitations: [state one limitation]. The 90% CI means that under repeated sampling, the true readiness would fall within this range 90% of the time." Full credit requires identifying both problems AND providing a corrected statement with CI and limitation.

SA-6. Full credit: any two WFF tracks correctly identified with an example — SL 4A (Intelligence): ORSA supports intel fusion by quantifying probability of enemy COAs or estimating adversary logistics capacity; SL 4B (Fires): ORSA supports fires by modeling target prioritization or computing probability of kill; SL 4C (Movement & Maneuver): ORSA supports M&M by optimizing route selection or modeling force employment timing; SL 4D (Sustainment): ORSA supports sustainment by forecasting fuel/ammunition demand or optimizing distribution; SL 4E (Protection): ORSA supports protection by modeling threat probability distributions or analyzing force protection trade-offs; SL 4F (Mission Command): ORSA supports MC by building dashboard analytics or quantifying COA risk for commander decision-making. Each response must identify the correct SL 4 letter (A–F) and provide a concrete product example for full credit.

---

*USAREUR-AF Operational Data Team TM-40G Post-Test | Version 1.0 | March 2026*