


DRAFT — UNOFFICIAL — NOT FOR OPERATIONAL USE

CONCEPTS GUIDE

# SL 5G



---

## CONCEPTS GUIDE — SL 5G COMPANION — ADVANCED OPERATIONS RESEARCH / SYSTEMS ANALYSIS (ORSA) — MAVEN SMART SYSTEM (MSS)

---

*Specialist Course Manual*

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

# CONCEPTS GUIDE — SL 5G COMPANION — ADVANCED OPERATIONS RESEARCH / SYSTEMS ANALYSIS (ORSA) — MAVEN SMART SYSTEM (MSS)

**Forward:** At SL 5G level, the ORSA is not just answering questions — they are shaping which questions get asked. **Purpose:** Extends the analytical mental models of the TM-40G Concepts Guide to advanced ORSA practice on MSS. Prerequisite: TM-40G Concepts Guide and SL 4G qualification. *HQ USAREUR-AF · v1.0 · 2026 · DISTRIB: USG only*

## SECTION 1 — FROM ANALYSIS TO ADVICE: THE ADVANCED ORSA'S SHIFT

### 1-1. The Core Transition

**BLUF:** At SL 5G level, the ORSA is not just answering questions — they are shaping which questions get asked.

A SL 4G ORSA operates in reactive mode by design: receive a requirement, scope the analysis, execute the method, deliver the product. That discipline is correct at SL 4G. Staying in lane builds accuracy and respect for the decision chain.

A SL 5G ORSA who operates only in reactive mode is failing. Senior leaders often know their operational problem but not which framing is analytically tractable, what data exists, or which methods produce reliable answers in their conditions. The SL 5G ORSA fills that gap. This is not overstepping — it is fulfilling the advisory function that distinguishes an ORSA from a data technician.

### 1-2. Reactive vs. Proactive Analytical Advising

Mode	Description	When Appropriate
Reactive	Answer the question as asked	Framing is analytically sound; data supports it
Reframing	Propose a modified question that is more tractable	Original question cannot be answered reliably with available data

Mode	Description	When Appropriate
Proactive	Surface an analytical question the commander has not asked	Your analysis reveals a risk or opportunity not yet on the decision radar

Most ORSAs underuse reframing and proactive advising because they feel like overreach. They are not — provided they go through the proper advisory channel and are paired with transparent analytical reasoning.

**Vignette — V Corps Readiness Analysis:** A G3 asks for a readiness trend analysis across assigned BCTs for the next 30 days. The ORSA runs the analysis and finds the trend is not the analytically interesting question — one BCT's equipment readiness is diverging from its personnel readiness in a pattern that predicts a near-term operational availability cliff. The ORSA's job: answer the trend question *and* surface the divergence with a clear explanatory note. That is proactive advising, not overreach. The commander decides what to do with it.

### 1-3. Developing Analytical Influence Without Overstepping

Analytical influence is built through a predictable pattern: be right more often than not, be honest about uncertainty, never oversell, never hide bad news. A single episode of overconfidence that later proves wrong damages credibility in a way that takes months to rebuild.

**Rule 1:** State the question you are answering, not just the result. Every product should begin with an explicit statement of the analytical question as scoped. If it was modified from the original request, note it.

**Rule 2:** Accompany every recommendation with the condition under which you would change it. "My recommendation is X. I would revise to Y if [specific data or event]." This signals maturity, builds trust, and gives leadership a clear trigger for returning to the analyst when conditions change.

## SECTION 2 — MULTI-METHOD ANALYSIS: WHEN AND HOW TO COMBINE METHODS

### 2-1. The Single-Method Temptation

**BLUF:** Most real operational problems require more than one analytical method. The risk is not in combining methods — it is in combining them poorly.

A force-generation readiness problem is simultaneously descriptive (where are we now?), diagnostic (why are specific units lower?), predictive (where will we be in 30 days?), and prescriptive (what maintenance allocation produces the best collective outcome?). A single trend analysis answers one layer. The SL 5G ORSA understands all four layers and can engage any on demand.

## 2-2. Designing a Multi-Method Analytical Campaign

---

An analytical campaign is a planned sequence of analyses addressing a problem set over time. Key design decisions:

**Sequencing:** Descriptive analysis typically precedes diagnostic, which precedes predictive, which precedes prescriptive. Operational time pressure often compresses this, but understanding the natural sequence helps when making informed shortcuts.

**Method selection:** Choose methods based on data available, decision time horizon, and required confidence — not sophistication. A linear regression with strong theoretical grounding is usually more defensible than a neural network with ambiguous feature importance. Prefer the simpler method unless the complex one demonstrably outperforms it out-of-sample.

**Handoff points:** Define explicitly where each method's output becomes the next method's input. These are where errors compound. Validate outputs at each stage before passing forward.

## 2-3. The Risk of Method Stacking

---

Method stacking is adding analytical complexity without adding insight — often driven by a desire to demonstrate rigor rather than serve the decision.

Signs of method stacking: - The product requires more than two minutes to explain the method before explaining the result - Uncertainty in each method's output compounds, making the final confidence interval meaningless - Methods produce conflicting results and the analyst averages them or picks the most favorable without a principled basis

When methods conflict, that conflict is itself analytically important. It often indicates the problem is not well-posed, the data has a structural issue, or the phenomenon is genuinely uncertain. Report the conflict explicitly.

**The one-clear-result principle:** When a simpler analysis produces a clear result and a more complex one produces an uncertain or conflicting result, the simpler result is usually more useful for decision-making. Document the complex analysis as a sensitivity check, not the primary product.

---

# SECTION 3 — WARGAMING AND SENSITIVITY ANALYSIS AS ORSA TOOLS

## 3-1. Analytical Wargaming on MSS

---

**BLUF:** Wargaming and sensitivity analysis are the advanced ORSA's primary tools for stress-testing models and communicating uncertainty to commanders.

Analytical wargaming in the ORSA context uses scenario modeling to stress-test analytical models and assumptions, not to explore adversary COAs. On MSS, this typically means constructing alternative parameter sets in Pipeline code and running the model under each.

**Three-scenario structure (standard minimum):**

Scenario	Parameter Setting
Base case	Central estimate of all uncertain parameters
Pessimistic	Parameters set to values producing the worst analytically plausible outcome
Optimistic	Parameters set to values producing the best analytically plausible outcome

This structure forces explicit documentation of the assumptions that drive the range of outcomes. That documentation is the analytical product.

**DEFENDER Exercise Vignette:** In preparing analytical support for a multi-division logistics simulation, the ORSA team runs three scenarios: base case assumes host-nation support at historical rates; pessimistic assumes 40% degradation (consistent with contingency planning assumptions); optimistic assumes marginal improvement. The pessimistic case reveals a Class V shortfall the base case conceals. A separate 10th AAMDC request for interceptor availability analysis under the same three scenarios reveals that the pessimistic case exhausts theater IAMD reserves 72 hours earlier than the base case — a finding with operational implications the original logistics-only analysis would not have surfaced.

### 3-2. Sensitivity Analysis — The Most Important Quality Check

**NOTE**

SL 4G now covers advanced multi-variable sensitivity analysis for MOE/MOP, including simultaneous variation of multiple parameters and interaction effects. Review that material before applying the single-variable methods below to multi-criteria assessments.

Sensitivity analysis answers: which assumption, if wrong, would most change the recommendation?

This is the SL 5G ORSA's most important quality check and is systematically underused. Most analysts run their model, check that results are sensible, and deliver. Sensitivity analysis requires deliberately breaking the model — varying one assumption at a time to the edge of plausible range.

**The threshold question:** For each key assumption, ask: at what value does the recommendation change from one COA to another? That threshold is the analytically important number.

**Presenting sensitivity results:** "Here is the recommendation under our most likely assumptions. Here is the single assumption that most influences it. If that assumption proves wrong in this direction, we should revisit." This is honest risk communication, not hedging. A commander who understands the key

assumption can provide better ISR collection guidance.

Sensitivity Presentation Approach	Appropriate When
Tornado diagram (ranked variable importance)	Many parameters; need to identify which matter most
Scenario table (3-5 scenarios)	Key uncertainties are discrete, not continuous
Threshold statement ("X must be below Y for this recommendation to hold")	Single dominant uncertainty; time-constrained brief
Full Monte Carlo CDF	Probability distributions are defensible; audience is analytically literate

## SECTION 4 — DECISION ANALYSIS UNDER UNCERTAINTY: THEATER-LEVEL APPLICATION

### 4-1. Frameworks for Decisions Under Uncertainty

**BLUF:** The appropriate decision-analysis framework depends on what is known about probabilities, what kind of error is more costly, and what the decision-maker's actual objective is. At theater level, these conditions are often ambiguous and must be established before analysis begins.

Framework	Best Suited For	When to Use
Multi-Criteria Decision Analysis (MCDA)	Decisions with multiple relevant criteria — readiness, risk, cost, coalition interoperability	Force-structure decisions, resource allocation, planning alternatives
Expected Utility	Decisions where probabilities can be reasonably estimated	Routine resource allocation; force generation
Minimax Regret / Robustness	Deep uncertainty; catastrophically asymmetric downside	Force-protection and posture decisions in USAREUR-AF
Satisficing	Meeting minimum acceptability thresholds across all criteria	Coalition planning where failure across any dimension is unacceptable

### 4-2. The USAREUR-AF Uncertainty Challenge

Three challenges specific to USAREUR-AF that most OR textbooks do not address:

**Multi-nation data incompatibility.** Coalition partner readiness data arrives in different formats, at different latencies, against different definitions of the same field. Treating it as a homogeneous dataset without harmonization is a significant analytical error.

**Compressed decision cycles.** Theater-level decisions in a crisis environment may have hours, not days, for analytical support. Design products that can be updated rapidly: modular analyses where the central estimate can be refreshed without rebuilding the entire pipeline.

**Political sensitivity of coalition analysis.** An MCDA product showing one partner nation's contribution as the weakest link is analytically correct and politically sensitive. Both are true. The senior ORSA navigates this by separating the analytical finding from the presentation venue. The analytical record must be complete; the decision brief is calibrated to audience and channel. This is appropriate handling of sensitive assessments, not dishonesty.

### 4-3. Avoiding Decision Paralysis

---

The most common failure when presenting uncertainty to commanders: inducing paralysis. The commander receives an honest picture of uncertainty and defers a decision that needs to be made.

The antidote is not hiding uncertainty — it is structuring the presentation so uncertainty is accompanied by a conditional recommendation: "Given what we know now, Option B dominates. If [specific intelligence] changes, Option C becomes preferable. Recommend Option B with a decision review trigger at [specific event or time]." This gives a clear action, a clear revisit condition, and confidence the analysis will update.

---

## SECTION 5 — ANALYTICAL PRODUCTS AT THE THEATER LEVEL

### 5-1. How Audience Changes the Product

---

**BLUF:** The underlying analysis does not change based on audience. The product — framing, depth, confidence language — changes significantly.

More senior audiences need analysis that is more clearly scoped, more honestly uncertain, and more decisionally relevant — not necessarily more methodologically complex. Briefing the optimization algorithm to the four-star is a category error.

## 5-2. Calibrating Depth to Decision Time Horizon

Decision Time Horizon	Product Must Provide	Omit
Hours (crisis action)	Single recommendation, key assumption, decision trigger	Method detail, full uncertainty characterization
Days (deliberate planning)	Range of alternatives, sensitivity on key variables, criteria weighting rationale	Derivation of all input data
Weeks (campaign assessment)	Full multi-scenario analysis, assumption register, pipeline documentation	Executive-level framing (covered separately)
Months (posture/force structure)	Complete analytical campaign, external review record, alternative framings	Tactical-level precision

An analyst who produces a 40-page methodology report for a 48-hour decision has misallocated effort. Calibrating analytical depth to decision time is a hallmark of senior-level ORSA maturity.

## 5-3. The Operations Assessment Working Group

### NOTE

SL 4G now addresses the ORSA's role in the operations assessment working group (OAWG), including structured assessment frameworks and coordination with the G3/G5. The theater-level product guidance below assumes familiarity with that OAWG foundation; advanced ORSAs are expected to lead or advise the OAWG analytical methodology, not just contribute products to it.

## 5-4. Coalition Visibility Considerations

Theater-level products in USAREUR-AF frequently have coalition visibility. Key considerations:

**Releasability.** Confirm releasability level before building the product. A pipeline that ingests NOFORN-flagged data cannot produce a RELIDO product without data scrubbing logic built in from the start. Ask before the pipeline is built.

**Partner nation data sovereignty.** Some NATO partner data can be used in analysis but not released back to other partners in disaggregated form. Aggregate visualizations that do not identify individual partner performance by name are often the solution.

**Assumption transparency with partners.** When assumptions are built on US intelligence estimates not shared with partners, the product may not be fully reproducible by partner analysts. Note this explicitly in methodology documentation.

## SECTION 6 — CROSS-DOMAIN DATA INTEGRATION

### 6-1. The Appeal and the Hazard

#### NOTE

SL 4G now covers structured MOE/MOP indicator frameworks and multi-variable sensitivity analysis that directly apply to cross-domain risk quantification. Ensure familiarity with the SL 4G assessment taxonomy (Section 1-7, Table 1-2) before designing cross-domain integration products at theater level.

**BLUF:** Cross-domain data integration is where the most analytically interesting theater-level insights live — and where the most analytically dangerous errors originate.

Personnel data, logistics data, maintenance data, and operational tempo data all describe the same force from different angles. Integration can reveal relationships invisible to any single domain — logistics data may predict operational readiness before maintenance data does. These relationships are real and analytically valuable. They are also dangerous when over-interpreted in small operational populations where spurious correlations are common.

### 6-2. Cross-Domain Risk Quantification

The starting question for any cross-domain integration: what is the minimum sample size to detect a relationship of operationally meaningful size with acceptable confidence? In theater operational analysis, the answer is frequently "more than we have." A correlation observed across 12 BCTs over six months is a weak signal, not a finding.

Three risk factors specific to cross-domain military analysis:

Risk Factor	Description	Mitigation
Definition inconsistency	Personnel "readiness" and logistics "readiness" use different formulas	Document and review all harmonization steps
Latency mismatch	Different domains have different reporting latencies; integrating without accounting for this produces false temporal correlations	Align to a common reporting window before analysis
Reverse causality	In small operational populations, shared drivers (e.g., high OPTEMPO) create coincident degradation across domains that looks like cross-domain prediction	Distinguish driver relationships from coincident degradation before making causal claims

### 6-3. NATO Exercise Coalition Data Integration — Vignette

During a large-force NATO exercise, an ORSA team integrates partner-nation reporting across five allied formations plus 10th AAMDC IAMD readiness data and SETAF-AF partner-force status feeds from the Africa AOR. Data arrives in six formats, against four readiness definitions, at three reporting frequencies.

**Correct approach:** 1. Produce a harmonized readiness index with documented transformation rules for each partner's data — note all assumptions 2. Report a confidence level for each partner's data based on format fidelity, reporting timeliness, and internal consistency 3. Present the coalition picture with partner-specific confidence annotations, not a single aggregate obscuring data quality variation 4. Flag the two harmonization assumptions that most affect the aggregate picture, with the directional impact of each

**Incorrect approach:** Average the available data without harmonization and present a clean coalition-wide number. That number is precise and meaningless.

## SECTION 7 — BUILDING AND SUSTAINING AN ANALYTICAL PIPELINE

### 7-1. The Rotation Problem

**BLUF:** An analytical product that only one person can run is a liability, not an asset.

USAREUR-AF ORSA billets turn over. A complex analytical pipeline built by one analyst over six months and left undocumented will be abandoned by their replacement — not out of laziness but out of rational triage. The intelligence community developed tradecraft documentation standards specifically to solve this problem: every significant analytical product maintains a methodology note sufficient for another analyst to reproduce the work without access to the original analyst. Apply this standard to operational ORSA on MSS.

### 7-2. Pipeline Transferability Requirements

Requirement	Standard
Input documentation	All data sources, field definitions, and transformation logic documented in-code or in an attached methodology note
Assumption register	All analytical assumptions listed with rationale and date last reviewed
Run instructions	A new analyst can run the pipeline end-to-end from written instructions without assistance
Output interpretation guide	What each output means, expected values, and values that should trigger investigation

Requirement	Standard
Known limitations	What the pipeline cannot do; data conditions causing failures; known edge cases
Update procedure	How to refresh with new data; which parameters may need recalibration over time

### 7-3. Automation Decision Framework

Automate when: - Transformation logic is stable and unlikely to change with operational context - The data source is reliable and the ingestion pathway is validated - The output is well-defined and the failure mode is observable (runs or errors — no silent failures) - A human review gate exists between the automated output and any action it triggers

Do not automate when: - Analytical logic requires contextual judgment varying with operational conditions - The data source has known quality variability requiring human assessment before processing - The pipeline output could be used to automate a decision that should retain human review

The most dangerous pipeline is one that runs reliably 95% of the time and fails silently in the remaining 5% — producing an output that looks correct but is not. Design for detectable failure over undetectable failure.

## SECTION 8 — ETHICS AND RESPONSIBILITY IN HIGH-STAKES ANALYSIS

### 8-1. The Weight of the Senior ORSA Role

**BLUF:** When analytical outputs inform force protection, personnel assessments, or coalition posture decisions, the professional obligation to accuracy is not abstract. It is personal and direct.

At battalion or brigade level, an ORSA error is consequential and correctable: the allocation recommendation was wrong, the analysis is revised, the decision is revisited. The feedback loop is short.

At theater level, the feedback loop is longer, the stakes are higher, and the correction cost may be significant. A readiness assessment that overstates coalition partner capability affects posture decisions, alliance assurance commitments, and contingency planning. The error may not surface for months, by which time decisions have been made and resources committed.

### 8-2. Presenting Honest Uncertainty to Leadership That Wants Confidence

The most ethically acute scenario in theater-level ORSA: leadership asks a question, the honest answer is that the data does not support a confident conclusion, and leadership wants confidence.

**Wrong responses:** - Produce a confident answer by selectively using supporting data - Bury uncertainty in methodology footnotes no one reads - Verbally agree that the analysis supports a conclusion it does not

**Correct response:** State the finding honestly and propose a path to the confidence level needed. "The current data supports a low-confidence assessment of X. To reach medium confidence, we need [specific data or time]. Here is what we can state with confidence now." This respects the operational requirement, states the limitation clearly, and offers a constructive path forward.

### 8-3. When to Decline to Produce an Analysis

---

Conditions under which the professional ORSA obligation is to decline: - The data does not support the confidence level the decision requires, and that gap cannot be closed in available time - The analytical question as framed contains a hidden assumption that predetermines a conclusion regardless of the data - The requested analysis would require using data outside its authorized use scope

Declining is done through the appropriate supervisory channel with a written explanation. The record protects the analyst and provides leadership the information needed to seek alternative analysis or adjust the requirement.

### 8-4. Personnel and Force Protection Analysis

---

When ORSA methods are applied to personnel data — identifying high-risk individuals, predicting retention, flagging adverse event patterns — the outputs have direct effects on careers and welfare.

**Rule:** If a person's career or safety could be affected by acting on your analytical output, that output requires higher confidence standards, explicit uncertainty communication, and documented review before it informs a decision.

---

## SECTION 9 — ADVANCED FAILURE MODES: WHAT SL 5G ORSAS GET WRONG

### 9-1. Overview

---

**BLUF:** Senior ORSAs make different mistakes than junior ORSAs. They are less likely to make technical errors and more likely to make judgment errors — anchoring, irrelevance, and communication failures that technical competence does not prevent.

## 9-2. Summary Table — Advanced Failure Modes

Failure Mode	Description	Diagnostic Question
Anchoring	Defending prior conclusions rather than re-examining them with new data	"Am I updating beliefs based on new data, or filtering data to match prior conclusions?"
Operational irrelevance	Technically correct analysis that does not map to an actionable decision	"What decision does this analysis support, and is it still pending?"
Communication failure	Sophistication that obscures the decision-relevant result	"Can a non-ORSA colleague state the recommendation in 30 seconds?"
Pipeline fragility	Analysis that only the original analyst can run	"Could a replacement analyst run this from documentation alone?"
Method stacking	Adding complexity without adding insight	"Does each method add to the recommendation, or just to the word count?"
Confidence overstating	Results presented with more certainty than the data supports	"Have I quantified and reported the key uncertainties explicitly?"

## 9-3. Anchoring to Prior Analysis

**NOTE**

SL 4G introduces Bayesian assessment updating as a formal method for incorporating new data into prior estimates. The anchoring failure described below is distinct from principled Bayesian updating — the former is cognitive bias; the latter is a disciplined framework for revising assessments. Understand the SL 4G Bayesian material before applying the diagnostic questions here.

The most common senior ORSA failure mode: rerunning a prior analysis with new data and treating prior conclusions as the baseline to defend rather than the hypothesis to test.

Anchoring manifests as: expecting conclusions similar to last cycle; investigating anomalies only when they produce significantly worse results; adjusting assumptions to restore prior conclusions rather than examining why new data produces different ones.

The distinction is between methodological consistency (correct) and conclusion anchoring (incorrect). Rerun every analysis as if seeing the data for the first time.

## 9-4. Analytically Correct but Operationally Irrelevant

A technically impeccable analysis that does not address a decision the commander can actually make is wasted effort. Check operational relevance twice: at the start (does this question map to a decision?) and before delivery (is that decision still pending?).

## 9-5. Analytical Sophistication Obscuring Communication

The senior ORSA can fail to communicate results effectively because they embed technical caveats and uncertainty quantification in ways that make the product difficult to act on. The test: give the briefing draft to a capable non-ORSA colleague and ask what decision the analysis supports. If they cannot answer in 30 seconds, the product needs revision — not the analysis, but the communication structure.

At theater level, a product that cannot be briefed in five minutes to a senior leader is rarely briefed. Both analytical quality and communicability are required.

## 9-6. Building Pipelines That Only the Original Analyst Can Maintain

The specific failure pattern: an ORSA builds a high-value pipeline over 3-6 months, becomes the recognized expert, and is then reassigned. The replacement cannot reproduce results due to undocumented environment dependencies, customized input file structures, or analytical logic never committed to documentation. Within 90 days, the pipeline is abandoned.

The fix is documentation discipline maintained throughout the build, not at the end. Write the methodology note as the pipeline is built. Assume you will be reassigned the day after you finish.

# APPENDIX A — CONCEPTS GUIDE SELF-ASSESSMENT

Before proceeding to SL 5G task procedures, confirm you can answer the following from memory:

1. What distinguishes reactive analysis from proactive analytical advising, and what conditions warrant each?
2. Name three methods appropriate for different layers of a multi-method analytical campaign. What determines which layer each addresses?
3. What is sensitivity analysis, and what is the "threshold question" it seeks to answer?
4. When is minimax regret more appropriate than expected utility for a theater-level decision? Give an operational example.
5. How does a product for a theater commander differ from the same underlying analysis delivered to a brigade commander?
6. What are the three primary failure modes in cross-domain data integration in small operational populations?
7. What six elements constitute a transferable analytical pipeline under SL 5G standards?
8. What is the correct professional response when leadership requests a confidence level the data cannot support?

9. Identify the most common senior-level failure mode and describe how to test whether you are exhibiting it.

## APPENDIX B — CROSS-REFERENCE TO SL 5G CHAPTERS

Concepts Guide Section	Corresponding TM-50G Chapter
Section 1 (From Analysis to Advice)	Chapter 1, Chapter 7
Section 2 (Multi-Method Analysis)	Chapter 2, Chapter 4
Section 3 (Wargaming and Sensitivity)	Chapter 5
Section 4 (Decision Analysis Under Uncertainty)	Chapter 6
Section 5 (Theater-Level Products; Operations Assessment WG)	Chapter 7
Section 6 (Cross-Domain Integration)	Chapter 4, Chapter 8
Section 7 (Building Analytical Pipelines)	Chapter 8
Section 8 (Ethics and Responsibility)	Chapter 7, Appendix B
Section 9 (Advanced Failure Modes)	All chapters — see Safety Summary

## APPENDIX C — PEER SL 5 CROSS-REFERENCES AND WFF INTEGRATION

**Peer SL 5 Publications.** The following advanced-track publications share overlapping concerns with SL 5G ORSA practice. Coordinate with practitioners in these tracks rather than operating in isolation.

Publication	Track	Coordination Point
SL 5H	Advanced AI Engineer	Productionizing ORSA models; evaluation methodology
SL 5M	Advanced ML Engineer	ML methods feeding ORSA analytical products
SL 5J	Advanced Program Manager	Portfolio-level OR program governance
SL 5K	Advanced Knowledge Manager	Capturing and surfacing ORSA-derived insights
SL 5L	Advanced Software Engineer	Platform infrastructure supporting ORSA pipelines
SL 5N	Advanced UI/UX Designer	Data visualization design; dashboard UX for ORSA outputs

Publication	Track	Coordination Point
SL 5O	Advanced Platform Engineer	Infrastructure for analytical workloads; compute scaling

**WFF Operational Consumer Note.** Advanced ORSA analytical work ultimately supports the six Warfighting Function (WFF) tracks: Intelligence (SL 4A), Fires (SL 4B), Movement and Maneuver (SL 4C), Sustainment (SL 4D), Protection (SL 4E), and Mission Command (SL 4F). These practitioners are the operational consumers of readiness analyses, logistics optimizations, campaign assessments, and decision support products. The judgment questions addressed in this Concepts Guide — which method fits the problem, how to present uncertainty, when to advise a different question — are ultimately answered in terms of what a WFF staff section needs to make a decision.

---

*CONCEPTS GUIDE — SL 5G COMPANION // ADVANCED ORSA HEADQUARTERS, UNITED STATES ARMY EUROPE AND AFRICA // WIESBADEN, GERMANY // 2026*

DRAFT