


DRAFT — UNOFFICIAL — NOT FOR OPERATIONAL USE

CONCEPTS GUIDE

SL 5K



CONCEPTS GUIDE — SL 5K: ADVANCED KNOWLEDGE MANAGER — 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 5K: ADVANCED KNOWLEDGE MANAGER — MAVEN SMART SYSTEM (MSS)

Forward: At SL 5K level, you are no longer maintaining a knowledge system. You are designing how the organization learns — a fundamentally different mandate with fundamentally different consequences if executed poorly. **Purpose:** Extends mental models from TM-40K Concepts Guide to advanced knowledge management on MSS. Prerequisite: TM-40K Concepts Guide and SL 4K qualification. *HQ USAREUR-AF · v1.0 · 2026 · DISTRIB: USG only*

SECTION 1 — FROM KNOWLEDGE MANAGER TO ORGANIZATIONAL LEARNING ARCHITECT

1-1. The Shift in Role

BLUF: At SL 5K level, you are no longer maintaining a knowledge system. You are designing how the organization learns — a fundamentally different mandate with fundamentally different consequences if executed poorly.

A SL 4K KM builds capture workflows, validates observations, structures repositories, and surfaces knowledge to users. Their scope is a unit. When they build poorly, that unit loses institutional memory at every rotation.

A SL 5K Advanced KM designs the infrastructure within which SL 4K KMs operate. Their scope is a theater — III Corps, V Corps, 21st TSC, 10th AAMDC, 56th MDC-E, SETAF-AF, USAREUR-AF HQ, and NATO partners. When they design poorly, the result is not a failed system — it is a theater of failed systems, each isolated, each redundant, collectively representing a massive organizational investment in knowledge that cannot be found, shared, or used.

1-2. From Capture to Learning Infrastructure

NOTE

SL 4K now codifies the FM 6-0 five-step KM process, three KM tasks, six KM tool categories, and the five integrating processes for enterprise KM across echelons. The theater-scale learning architecture described below builds on that doctrinal foundation; review TM-40K Section 1-5b before designing cross-echelon KM systems.

Tactical KM (SL 4K) is primarily concerned with capture: getting observations into a system before knowledge walks out the door at RIP/TOA. Strategic KM (SL 5K) is concerned with the full learning cycle:

Phase	Tactical KM Focus	Strategic KM Focus
Capture	Build submission workflows	Design taxonomy that makes submissions findable at scale
Validation	Review individual lessons	Design the quality tier architecture and human review pipeline
Distribution	Surface lessons to unit users	Design cross-domain linkages so logistics lessons reach planners
Application	Notify users of relevant lessons	Instrument the system to detect whether recommendations are used
Doctrine	Escalate high-value lessons	Design the pathway from validated lesson to formal guidance
Retirement	Archive superseded material	Design governance cycles that systematically retire outdated knowledge

The hardest phase is not capture — it is application. An organization that captures ten thousand lessons and applies none of them has not learned. It has archived. The Advanced KM designs for application as the primary outcome, with capture as a necessary input.

1-3. The Compounding Learning Organization

A compounding learning organization has three properties:

Lessons propagate. A lesson from V Corps during Saber Strike reaches a 21st TSC planner before the next exercise. This requires cross-domain linkage, a distribution mechanism, and a taxonomy that makes the lesson findable by someone who did not submit it.

Doctrine incorporates experience. Validated lessons feed into updated SOPs, TTPs, and training guidance. This requires a formal pathway from the KM system to the publication authority — and a publication authority that actually consumes KM outputs.

Rotation cycles do not reset the system. An organization that rebuilds its knowledge system every 18 months because the KM who built it PCS'd has not compounded anything — it has run in place. Compounding requires designed durability: documentation, succession planning, and system architecture that does not depend on a single person's expertise.

1-4. Vignette — The Defender Europe Lessons Learned Problem

During Defender Europe 22, V Corps collected 2,400+ observations. The V Corps KM team validated 847 lessons; 312 were categorized as high-value warranting cross-domain or higher-echelon review.

At exercise conclusion, 309 of 312 high-value lessons remained within the V Corps knowledge system. Three were formally escalated to USAREUR-AF doctrine staff. No distribution mechanism routed logistics lessons to 21st TSC. No IAMD coordination lessons reached 10th AAMDC. No multi-domain effects observations reached 56th MDC-E. No linkage existed between operational planning lessons and the JOPP training program at Grafenwoehr. No coalition-releasable version of any lesson was produced.

Defender Europe 23 planning began 14 months later. Planners had no structured access to Defender 22 lessons — the information existed but was not findable, not linked to current planning questions, and not available to coalition partners.

This is not a data problem. It is a KM architecture problem. The Advanced KM's role is to design the system so that Defender Europe 23 planners can, within minutes, surface every relevant logistics, communications, and fires lesson from Defender 22 — regardless of which staff section captured it, regardless of security level, and regardless of whether the original submitter is still assigned to USAREUR-AF.

SECTION 2 — ENTERPRISE KNOWLEDGE ARCHITECTURE

2-1. The Scale Problem

BLUF: A knowledge architecture that works for a battalion breaks at Corps. A knowledge architecture that works for a single functional domain breaks when cross-domain requirements are added. Enterprise KM architecture must be designed for scale from the beginning — retrofitting is expensive and frequently fails.

The USAREUR-AF theater presents a specific challenge: multiple echelons (Army, Corps, Division, Brigade), multiple functional domains (G2, G3, G4, G6, Civil Affairs, Aviation, Fires, Sustainment, Medical, Legal, CEMA, SOF, multi-domain), multiple classification levels, and coalition partner integration — all served by knowledge systems consistent enough to share information but flexible enough to meet distinct organizational requirements.

Centralized architecture fails predictably: no single taxonomy satisfies all functional domains, centralized review bottlenecks at scale, units stop submitting when their requirements are not met, and the system becomes dominated by the most aggressive functional area. Full decentralization produces knowledge silos requiring manual cross-unit coordination that rarely happens. The correct approach is federated architecture.

2-2. Federated vs. Centralized — The Design Choice

Characteristic	Centralized	Federated	Recommended for USAREUR-AF
Taxonomy control	Single authority	Distributed with shared core	Federated with mandatory core
Review/validation	Central KM team	Unit KMs with central coordination	Federated
Cross-domain discovery	Easy (same system)	Requires linkage design	Design linkages explicitly
Unit autonomy	Low	High	High, with governance guardrails
Coalition integration	Single interface point	Multiple interface points	Centralized interface, federated storage
Resilience to personnel loss	Low (single point of failure)	High	Federated
Consistency of quality	High (if central team is strong)	Variable	Requires quality tier standards

The USAREUR-AF architecture should be federated at the storage and curation level, with a mandatory shared taxonomy core enabling cross-domain and cross-echelon search. Each major formation (III Corps, V Corps, 21st TSC, 10th AAMDC, 56th MDC-E, SETAF-AF, USAREUR-AF HQ) maintains its own MSS knowledge repository. A Theater Knowledge Index aggregates metadata — not content — from all repositories, enabling theater-wide search without centralizing content or review authority.

2-3. Designing the Shared Taxonomy Core

The shared taxonomy core is the minimum set of classification dimensions every knowledge product in the theater must carry. It enables cross-domain search without requiring all units to adopt the same internal taxonomy.

Dimension	Required Values	Notes
Functional domain	G1, G2, G3, G4, G6, Civil Affairs, Aviation, Fires, Sustainment, Medical, Legal, CEMA, SOF, Multi-domain	Units may add sub-domains; must map to at least one core domain
Echelon	Theater, Corps, Division, Brigade, Battalion	Echelon at which the lesson is applicable
Exercise / operation	Defender Europe, Saber Strike, Combined Resolve, Swift Response, [named operation], Steady-state	Source event or context
Quality tier	Observation, Reviewed, Validated, Doctrine	See Section 4
Releasability	US-only, NATO-releasable, PfP-releasable, Public	See Section 6
Status	Active, Under review, Superseded, Retired	See Section 8

Units may extend their internal taxonomies as needed. They may not omit core dimensions — the Theater Knowledge Index depends on consistent core classification for cross-unit search.

2-4. The Theater Knowledge Index on MSS

The Theater Knowledge Index is a metadata aggregator — an MSS dataset containing core classification fields, a title, a summary, a URL back to the source repository, and the access control level for each knowledge product across all federated repositories.

Users search the index to discover relevant knowledge theater-wide. When they find a relevant item, the index points them to the originating unit's repository for full content access (subject to access control). The KM managing the Theater Knowledge Index is responsible for taxonomy consistency and index freshness — not for content ownership.

Critical property: no unit's knowledge is replicated into a central system it does not control. Each unit remains the authoritative owner of its own knowledge products. The theater-level system provides discovery without requiring custody transfer.

2-5. Vignette — V Corps and 21st TSC: One Theater, Two Architectures

V Corps operates a KM system centered on operational lessons: fires synchronization, ISR allocation, command post operations, information operations. Taxonomy is G-staff focused, exercise-driven, integrated with the JOPP review cycle.

21st TSC operates a KM system centered on sustainment lessons: distribution network design, maintenance backlogs, medical logistics, contractor support integration. Taxonomy is functional-area focused, tied to logistics readiness cycles, integrated with LOGSTAT reporting.

10th AAMDC operates a KM system centered on IAMD lessons: sensor-shooter integration, early warning procedures, missile defense coordination with NATO BMD architecture. 56th MDC-E captures multi-domain effects lessons: EW/cyber/space integration, cross-domain fires deconfliction, and joint targeting in contested electromagnetic environments. SETAF-AF maintains lessons focused on expeditionary operations in the Africa AOR: partner force coordination, austere logistics, and security cooperation engagement outcomes. Each system uses its own taxonomy, its own storage, and its own review cycle.

A lesson about fuel distribution during Defender Europe 22 — captured in 21st TSC's system under "POL distribution, forward" — is directly relevant to V Corps planners designing fuel support for Defender Europe 23. Without a shared taxonomy mapping 21st TSC's "POL distribution" tag to the theater-level "Sustainment" domain, and without a Theater Knowledge Index making 21st TSC lessons visible to V Corps search queries, that lesson never reaches the people who need it.

The shared taxonomy core and the Theater Knowledge Index together solve this problem without requiring either formation to restructure its internal KM system.

SECTION 3 — AI-ASSISTED KNOWLEDGE MANAGEMENT — THE MSS OPPORTUNITY

3-1. The Role of AI in KM Workflows

BLUF: AIP Logic and AI tools on MSS can substantially reduce the manual workload of knowledge management at scale. They can also introduce systematic quality failures if deployed without appropriate human oversight. Design principle: AI augments the KM workflow; it does not replace the human judgment required to validate and approve knowledge products.

At theater scale, a Corps-level exercise may generate 200–400 observations in 72 hours. A theater-level program running concurrent exercises may generate that volume daily. Manual review at that rate is not feasible without either a large dedicated KM staff or significant review latency. AI augmentation addresses this by handling high-volume, low-judgment portions of the workflow — leaving human reviewers focused on high-judgment work: validation, approval, and doctrine escalation decisions.

3-2. AI Augmentation Functions in KM Workflows

Function	What AI Does	Human Role	Risk if Human Role Is Removed
Auto-tagging	Suggests taxonomy classifications based on submission content	Review and confirm; correct systematic errors	Systematic misclassification; broken cross-domain discovery

Function	What AI Does	Human Role	Risk if Human Role Is Removed
Duplicate detection	Flags submissions similar to existing knowledge products	Review flagged pairs; decide whether to merge, link, or keep separate	Duplicate accumulation or inappropriate merging of distinct lessons
Pattern detection	Identifies recurring themes across large observation sets	Interpret patterns; decide whether pattern represents a systemic issue warranting escalation	False patterns; real patterns missed because AI framing is accepted without interrogation
Precedent recommendation	Surfaces related validated lessons when a new submission arrives	Review recommendations; decide whether to add cross-references	Irrelevant links that reduce signal-to-noise; relevant links missed
Summary generation	Produces plain-language summaries of complex technical lessons	Review and edit; ensure accuracy and appropriate classification	Inaccurate summaries distributed as authoritative; classification errors
Draft SOP/TTP generation	Produces structured draft guidance from validated lesson sets	Full review, edit, and approval by qualified SME and appropriate authority	AI-generated doctrine published as authoritative; errors propagate into training and operations

AI handles first-pass processing; humans handle all approval decisions. The speed benefit comes from reducing cognitive load on human reviewers — they adjudicate AI proposals rather than build from scratch, which is substantially faster without sacrificing judgment.

3-3. Designing the AI-Augmented Review Queue

Submitted observations enter the queue. An AIP Logic pipeline processes each submission and adds: suggested taxonomy tags with confidence scores, duplicate detection flags, linked related lessons (top 3–5 by relevance score), and a draft one-sentence summary.

The human reviewer sees the submission plus AI-generated metadata. They confirm or modify taxonomy, resolve duplicate flags, accept or reject related-lesson links, and edit the summary. Review time drops from 15–20 minutes (manual) to 3–5 minutes (AI-augmented). At Corps scale, this is the difference between a team of 6 KMs and a team of 2 managing the review queue.

NOTE

The confidence score on AI-suggested tags is a useful signal, not a threshold for auto-approval. Do not design workflows where high-confidence AI tags are automatically accepted without human review. The cost of systematic misclassification at theater scale — thousands of improperly tagged lessons — exceeds the labor cost of review.

3-4. Pattern Detection as an Intelligence Function

Pattern detection across large observation sets is qualitatively different from other AI functions — it operates on the corpus, not on individual submissions. An AI pipeline that analyzes 2,000 observations from a year of exercises and identifies "fuel distribution failures at forward support companies" as a recurring theme in 14% of logistics submissions is surfacing genuine organizational intelligence.

The Advanced KM designs for this function deliberately: maintain corpus quality (garbage in, garbage out), design queries that surface recurring themes rather than just individual keywords, and establish a governance process for acting on identified patterns — who reviews the pattern report, who decides whether to escalate to doctrine, who tracks whether recommended changes were made.

Pattern detection results must never be auto-published. A recurring theme in AI analysis may reflect a genuine systemic problem, a systematic bias in who submits observations, a classification error pattern, or an artifact of how the AI processes language. Human analysis is required to distinguish among these before any pattern-based recommendation is issued.

SECTION 4 — KNOWLEDGE QUALITY ARCHITECTURE

4-1. The Quality Problem at Scale

BLUF: At unit scale, a KM can maintain quality through direct involvement in every submission. At theater scale, direct involvement is impossible — quality must be designed into the system architecture.

A theater-scale system that treats a rough field observation with the same authority as a doctrine-reviewed SOP will be ignored. Users develop trust through experience — if early searches surface authoritative, useful results, users return. If searches surface unvalidated, contradictory, or outdated material, users stop searching and the system becomes an unused archive.

Quality architecture solves this by making quality visible and searchable. Users filter by quality tier, calibrate trust accordingly, and know that Validated lessons have been reviewed by a qualified SME with appropriate authority.

4-2. The Quality Tier Model

Tier	Label	Definition	Required Approver	Search Priority
Tier 0	Observation	Raw submission; not yet reviewed	None (submitter only)	Lowest; filtered out by default for general search
Tier 1	Reviewed	Checked for completeness, accuracy, and relevance by an SME	Unit-level SME	Available in search; displayed with tier label
Tier 2	Validated	Approved by a responsible authority	O-5 or above, or designated program authority	Prioritized in search results
Tier 3	Doctrine	Incorporated into formal guidance (SOP, TTP, EXORD annex, published doctrine)	Publication authority (varies by document type)	Highest priority; distinguished by tier in all search results

The quality tier is a field in the core taxonomy and a search filter in every Workshop application. Default search view for operational users surfaces Tier 2 and above. KMs and SMEs access Tier 0 and 1 for review work. All tiers are available to Advanced KMs managing the pipeline.

4-3. The Bottleneck Problem

Quality review is a bottleneck by design — human judgment cannot scale linearly with submission volume. The Advanced KM's job is not to eliminate the bottleneck but to design for throughput without sacrificing integrity. Five design levers:

Lever	Application
Distributed review authority	Functional domain SMEs validate within their domain. The KM coordinates the pipeline; does not validate all content personally.
Tiered review thresholds	Reserve full Tier 2 validation for lessons likely to be widely shared, cited in plans, or escalated to doctrine. Tier 1 is sufficient for many use cases.
AI pre-processing	Reduces per-submission review time substantially; multiplies throughput of a fixed reviewer population.
Review cycle batching	For non-time-sensitive submissions, batch by functional domain on weekly or bi-weekly cycles. Reduces context-switching cost for collateral-duty SME reviewers.
Queue visibility	Build dashboards showing queue depth, average review time by domain, and submissions aging past threshold. Makes the bottleneck visible; prerequisite for leadership resourcing decisions.

4-4. Integrity Without Bottleneck — The Advanced KM's Balance

The tension between throughput and integrity is managed, not resolved. Calibrate the system to operational tempo: tighter review gates during steady-state, faster throughput with appropriate tier labeling during high-tempo exercises.

The integrity risk is not a planner making an informed choice to use Tier 1 material. The risk is a system where quality tiers are inconsistently applied, where users cannot trust that a Tier 2 label means what it says, or where AI pre-processing has effectively replaced human validation without the design explicitly acknowledging that fact. Those are failure modes, not features.

SECTION 5 — CROSS-DOMAIN KNOWLEDGE INTEGRATION

5-1. The Silo Problem

BLUF: Knowledge silos are the default outcome of any enterprise KM system that does not explicitly design for cross-domain integration. Every staff section with a SharePoint folder, every unit with its own Foundry dataset, every functional area with its own taxonomy is, in isolation, building a knowledge silo. The Advanced KM designs the linkages that make cross-domain knowledge discoverable.

The operational cost of knowledge silos is easy to recognize in retrospect: a sustainment lesson about fuel distribution failure modes during cold-weather operations exists in the 21st TSC KM system, while a V Corps operational planner designs a cold-weather operation with a fuel support configuration that replicates the same failure mode — because they had no way to know the lesson existed. An IAMD coordination lesson captured by 10th AAMDC during a Patriot repositioning exercise is directly relevant to 56th MDC-E planners integrating air defense into a multi-domain task force — but it sits in a separate system with no cross-referencing mechanism.

Cross-domain integration does not require merging repositories, adopting a unified taxonomy, or centralizing review authority. It requires three design elements: shared classification dimensions, linked Object Types in the Ontology, and cross-domain search access.

5-2. Designing for Cross-Domain Discoverability

The shared taxonomy core (Section 2) provides the first element. When every knowledge product carries a functional domain tag from a shared controlled vocabulary, a search for "Sustainment" lessons surfaces relevant products from every repository using the shared core — regardless of each unit's internal taxonomy.

The Ontology provides the second element. In MSS, Object Types can be linked: a Lesson Object Type can have "related lessons" and "source exercise" links. When a G4 planner views a sustainment lesson, a properly linked Ontology shows related operational planning lessons from the G3 domain — because a KM or AI pipeline has explicitly created those relationships.

The third element is access. Cross-domain search is only useful if the user has read access to relevant repositories. Design access control to be permissive for search (broad read access to metadata and summaries) and appropriately controlled for full content (tiered by classification and need-to-know). The Theater Knowledge Index model (Section 2) implements this: index metadata is broadly accessible, full content access is gated at the originating repository level.

5-3. The Cross-Domain Link as a KM Work Product

Explicit cross-domain links do not create themselves. At theater scale, the Advanced KM cannot personally create every cross-domain link. The design solution is distributed cross-linking: when a KM validates a lesson in their domain, they add cross-domain links to related lessons in other domains. AI-suggested related lessons (from the augmented review queue, Section 3) reduce discovery work — the KM reviews suggestions and confirms or rejects them. The result is a progressively richer cross-domain link structure reflecting the judgment of distributed functional SMEs.

5-4. Vignette — The Fires-Sustainment Link That Prevented a Planning Error

During planning for a V Corps combined arms live fire, the G3 fires planner queried the theater knowledge system for lessons on artillery ammunition resupply during high-tempo fires operations. The system returned eight results — six from Fires domain sources and two from Sustainment domain sources.

The two Sustainment results — cross-linked by a 21st TSC KM who identified the relevance during validation — contained a Tier 2 validated lesson on ammunition convoy throughput constraints during contested road network conditions. The lesson was specific: 40% lower throughput than planning factors assumed during Defender Europe 22. The fires planner used those constraints to adjust planned fires tempo, requiring a larger pre-positioned ammunition stockpile. The sustainment requirement was identified during planning, not during execution.

Without the cross-domain link, the planner would have had no way to know the 21st TSC lesson existed. The cross-domain link was the single design element that converted a captured lesson into a planning decision.

SECTION 6 — COALITION KNOWLEDGE MANAGEMENT

6-1. The Coalition KM Challenge

BLUF: USAREUR-AF does not operate alone. Coalition KM requires designing a system that serves both US-only and coalition-releasable knowledge without creating two separate systems.

The default response to coalition knowledge sharing is a separate coalition system — a parallel repository with a separate taxonomy, separate review workflow, and separate user population. This doubles KM workload, creates divergent knowledge bases that are difficult to reconcile, and produces coalition partners with access to a curated fraction of theater knowledge rather than a genuinely useful operational resource.

The alternative is a releasability-tiered single system: one knowledge repository with a releasability tag on every knowledge product controlling who can access it. The KM creates one knowledge product; the system handles access control.

6-2. The Releasability Tag as a Design Element

The releasability tag is part of the mandatory shared taxonomy core (Section 2) and enforced at the Foundry dataset and Workshop application level through MSS access controls.

Releasability Level	Accessible To	Default for New Submissions	Modification Authority
US-only	US personnel with appropriate clearance	Yes — default for all new submissions	Originating unit KM, with OPSEC review for downgrade
NATO-releasable	All NATO member nation forces	No — requires explicit designation	Originating unit KM, with OPSEC and classification review
PfP-releasable	Partnership for Peace nation forces	No — requires explicit designation	Originating unit KM, with OPSEC, classification, and legal review
Public	Unrestricted	No — requires explicit designation and public affairs review	Designated publication authority only

WARNING

Downgrading releasability — particularly to NATO-releasable or PfP-releasable — is not a routine KM administrative action. It requires OPSEC review to ensure that the aggregated knowledge product does not reveal sensitive capability data, operational patterns, or personnel information even if individual elements are individually unclassified. Design content sanitization as part of the coalition-releasable designation workflow.

6-3. STANAG 4778 and NATO Information Management Policy

Coalition knowledge sharing operates under a governance framework the Advanced KM must design within. STANAG 4778 governs NATO lessons learned interoperability. NATO Information Management Policy governs what information can be shared across the Alliance and through what channels. Key practical constraints:

Constraint	Application
Information sharing agreements must precede technical access	A partner nation cannot be given access to coalition-releasable content on MSS without a signed information sharing agreement on file. The KM does not manage agreements — but must confirm they exist before granting access.
Coalition knowledge products must be explicitly designated	No knowledge product is coalition-releasable by default. US-only default ensures deliberate release decisions — protecting against inadvertent disclosure through system misconfiguration.
Architecture must support auditability	Coalition sharing requires demonstrating after the fact what was shared with whom and when. Foundry's access control and audit logging support this, but the KM must design the system to preserve those logs.

6-4. Designing for Coalition Without Creating Two Systems

Design Decision	Implementation
Separate Workshop views by releasability	US personnel see the full system. Coalition partner users see a Workshop application filtered to their releasability level. Underlying data is the same; view is access-controlled. Requires designing Workshop applications from the beginning to support releasability filtering as a core parameter.
Coalition-appropriate metadata	A knowledge product designated as NATO-releasable may require a sanitized title and summary not referencing US-only context. The KM workflow should include a review-and-edit step for coalition-appropriate language when designating a product as coalition-releasable.
Coalition feedback	Coalition partner submissions carry a source nation tag and the submitting nation's releasability designation. A lesson from a German Bundeswehr KM is not automatically

Design Decision	Implementation
integration	releasable to all NATO partners. The system must support multi-nation releasability tags, not a US-centric model.

SECTION 7 — MEASURING KNOWLEDGE SYSTEM EFFECTIVENESS

7-1. Beyond Count Metrics

BLUF: Submission volume, repository size, and taxonomy completeness are process metrics. They measure KM system activity, not effectiveness. An Advanced KM who reports "2,400 observations submitted" without reporting whether those observations were found and used has measured the input, not the outcome.

Process metrics are necessary for managing operations and identifying bottlenecks — but insufficient for demonstrating operational value. The Advanced KM needs a measurement architecture connecting KM system activity to operational outcomes.

7-2. The Effectiveness Metric Framework

Metric Type	What It Measures	How to Instrument on MSS	Operational Interpretation
Submission volume	Observation capture rate	Count of Observation objects by time period, domain, unit	Baseline activity; high volume does not indicate effectiveness
Review throughput	Quality tier advancement rate	Count of tier transitions (Obs → Reviewed → Validated) by time period	KM pipeline health; persistent backlogs indicate resource constraint
Search utilization	Whether users are querying the system	Workshop analytics: search query counts, unique users by week	Adoption rate; zero search means zero utility regardless of content quality
Search result engagement	Whether users are opening results	Workshop analytics: result click-through rate by tier and domain	Content quality signal; low engagement after search suggests results are not relevant
Cross-domain discovery	Whether cross-domain links are being followed	Workshop analytics: navigation events from one domain to another	Cross-domain integration effectiveness

Metric Type	What It Measures	How to Instrument on MSS	Operational Interpretation
Lesson citation	Whether lessons are cited in plans or products	Reference tracking: links from plan documents to validated lessons	Highest-value outcome metric; requires a citing convention and compliance
Repeated failure detection	Whether the organization is making the same mistakes twice	Periodic analysis: do new observations match themes of validated lessons in the same domain	Definitive organizational learning metric; requires periodic human analysis
Rotation continuity	Whether system usage continues through RIP/TOA	Search volume and submission rates during and after unit transitions	System durability metric; drops at transition indicate person-dependent systems

7-3. Foundry Workshop Analytics

MSS Foundry Workshop includes native usage analytics providing search query volumes, user engagement, and navigation patterns. Configure a KM Effectiveness Dashboard surfacing the key metrics above on a weekly or monthly review cycle.

Analytics data is most useful in context. A drop in search utilization the week after a major exercise ends is expected. The same drop three months into steady-state operations — with no corresponding decrease in submission volume — suggests users submitted but did not engage to retrieve. That pattern warrants investigation: the supply side is functioning, but the demand side is not.

7-4. The Repeated Failure Metric

The most operationally significant metric — whether the organization is making the same mistakes twice — is also the most difficult to instrument because it requires human interpretation.

Practical implementation: a quarterly analysis cycle in which the Advanced KM (or designated theater KM analyst) reviews newly validated lessons against the validated lessons corpus, identifies thematic matches, and produces a brief report: "These five new validated lessons appear to represent recurrences of previously validated lessons in the same domain. Recommended action: review whether prior lessons were incorporated into planning, training, or doctrine, and if not, escalate for action."

AI-assisted similarity detection can flag candidate matches for human review, but the judgment — is this the same failure? — is a human judgment. The value of the metric depends on that judgment, which is why it belongs at the Advanced KM level.

SECTION 8 — KNOWLEDGE LIFECYCLE MANAGEMENT AT SCALE

8-1. The Accumulation Problem

BLUF: A knowledge system that never retires outdated material will eventually become an obstacle to knowledge access. Users who cannot distinguish current lessons from superseded ones will stop trusting search results — and the system will become a liability rather than an asset.

At theater scale, no individual KM can maintain familiarity with the full repository. Currency management must be designed into the governance architecture. A theater KM system operating for five years without systematic retirement processes contains knowledge products from organizational structures that no longer exist, equipment that has been replaced, techniques superseded by updated doctrine, and observations thoroughly analyzed and incorporated into current guidance. That material is not wrong — it may be historically interesting — but its presence in active search results reduces signal-to-noise for users seeking current applicable guidance.

8-2. Knowledge Lifecycle Stages

Stage	Status Tag	Description	Trigger	KM Action
Active	Active	Current, applicable knowledge	Default on validation	None — maintain
Under review	Under review	Flagged for currency or accuracy review	Scheduled review cycle, or flagged by user or AI	Assign to reviewing SME
Superseded	Superseded	Replaced by newer knowledge product or updated doctrine	Reviewer determination	Link to superseding product; remove from default search
Retired	Retired	No longer applicable; archived for historical reference	Reviewer determination after supersedure period	Move to archive tier; accessible to KMs, not in operational search

Active and Under review products appear in operational search; Superseded and Retired products do not (but remain accessible to users who explicitly filter for them or navigate to the archive).

8-3. The Scheduled Review Cycle

Review Type	Frequency	Scope	Responsible Party
New validation review	Continuous	Tier 1 observations ready for Tier 2 validation	Domain SMEs per KM assignment
Annual currency review	Annual	All Tier 2 validated lessons older than 24 months	Theater KM, with domain SME review
Post-doctrine update review	Event-driven	All lessons in domains affected by a doctrine publication change	Theater KM; triggered by doctrine publication events
Post-equipment change review	Event-driven	All lessons referencing equipment undergoing fielding or replacement	Theater KM; triggered by G4 equipment transition notifications
Post-RIP/TOA review	Event-driven	All lessons associated with the transitioning unit	Outgoing unit KM, with theater KM oversight
Quarterly pattern review	Quarterly	New validated lessons vs. existing corpus for recurrence analysis	Theater KM or designated analyst

The post-RIP/TOA review is the highest-risk event for knowledge currency. When a unit transitions, the organizational context giving meaning to many of their knowledge products transitions with them. Unit-specific references, named individuals, internal system names, and exercise-specific context may no longer be interpretable by the incoming unit. The lifecycle review at RIP/TOA should determine whether unit-specific knowledge products are still meaningful in the incoming unit's context, or should be generalized, superseded, or retired.

8-4. Designing for Three Rotation Cycles

A durable KM system functions well after three RIP/TOA cycles (approximately 4–5 years) without significant architectural intervention. Person-dependent systems fail predictably at transition:

- System operation requires knowledge that lives primarily in the current KM's head rather than in documented procedures
- Taxonomy decisions were made informally and exist as institutional knowledge rather than written governance
- Cross-domain links were created by one person and are not understood by successors
- AI pipeline configurations are undocumented and require the original architect to maintain
- Review authority relationships are personal rather than structural (e.g., "ask MAJ Smith" rather than "submit to G3 KM branch chief")

Design explicitly against these failure modes:

- Every governance decision documented as a written policy or architecture decision record (ADR)

- Every taxonomy element with a written definition in the shared governance documentation
- Every AI pipeline with a runbook: what it does, how configured, how to modify, who to contact when it fails
- Every review authority relationship tied to a position, not a person
- Every system with a "keys in hand" handover package: current as of the last rotation, updated as part of every RIP/TOA process

The test: if the current Advanced KM were reassigned tomorrow, could their replacement understand and maintain the system from documentation alone? If the answer is no, the system is not yet durable.

SECTION 9 — ADVANCED FAILURE MODES — WHAT SL 5K KMS GET WRONG

BLUF: Advanced KM failures are more subtle than SL 4K failures. Sophisticated architecture can obscure fundamental problems until the system is deployed at scale, at which point remediation is expensive. A SL 4K KM who builds a poor system affects a unit. A SL 5K KM who builds a poor system affects a theater.

Failure Mode	Indicators	Correction
Architecturally impressive, operationally unused system	Submission forms with 15+ fields; search interfaces requiring taxonomy knowledge; knowledge products written for a KM audience, not an operational one; adoption metrics never built	Design for the end user first. If the people who need the system are not using it, the system has failed — regardless of technical quality.
Designing for capture without designing for retrieval	More design effort on the submission form than the search interface; no analytics on search utilization; users who know lessons exist but cannot find them; outdated material surfaced ahead of current material	Spend at least as much design effort on retrieval as capture. Test the search experience with operational users before deployment. Instrument for retrieval effectiveness from day one.
AI augmentation creating false confidence in knowledge quality	AI tag confirmation rates above 95%; systematic misclassification patterns persisting across the validation pipeline; reviewers who cannot articulate why a submission was classified as it was; AI-generated summaries published without material human editing	Design the review workflow to require independent human evaluation before displaying AI suggestions, or include a mandatory "disagree" step. Regularly audit AI suggestion accuracy against independent human classification. Train reviewers to interrogate AI proposals, not ratify them.
Failing to sustain	System health depends on a specific individual's sustained attention; when that	Design for minimum viable operation. Identify the absolute minimum KM resource level at

Failure Mode	Indicators	Correction
organizational commitment through rotation cycles	person or priority changes, submissions fall behind, review backlogs accumulate, and operational users who encounter stale results stop searching	which the system maintains acceptable quality. Brief it to leadership as a sustainment requirement. A system needing one skilled KM and four hours per week to maintain is one leadership can commit to.
Building knowledge systems without decommission plans	System outlives operational utility; users uncertain which system is authoritative; migration costs exceed original build costs; governance conflicts between old and new system custodians	Document the decommission plan at design time: when would this system be retired, what are the trigger conditions, what does data migration look like, who owns the archive. Decommission is a managed process, not an emergency.

SUMMARY — KEY PRINCIPLES FOR THE ADVANCED KM

Principle	Application
Design for application, not capture	The measure of a knowledge system is whether users make better decisions because of it. Capture is a means, not an end.
Federated with a shared core	Theater-scale KM works best when each organization owns its knowledge and a shared taxonomy enables cross-domain discovery. Enforce the core; leave the rest flexible.
AI augments; humans decide	Every AI function in a KM workflow has a corresponding human approval step. Speed comes from reducing cognitive load, not from removing human judgment.
Quality must be designed, not assumed	Build quality tiers into the architecture. Make quality visible to users. Design the review pipeline for throughput, not just rigor.
Cross-domain links are a work product	Discovery across functional domains does not happen automatically. Design the workflow to distribute cross-linking work across functional SMEs.
Coalition is one system, not two	The releasability tag enables coalition sharing without duplicating effort. Design for it from the beginning.
Measure outcomes, not activity	Submission counts measure effort. Lesson citation rates, repeated failure detection, and rotation continuity measure effectiveness.
Retire actively	An active retirement program is as important as an active capture program. A repository that never retires material degrades in signal-to-noise until it becomes unusable.

Principle	Application
Design for your third successor	Document everything. Tie authority to positions, not people. If the system requires you personally to function, it will fail at your next assignment.

PEER SL 5 CROSS-REFERENCES AND WFF INTEGRATION

Peer SL 5 Publications. Advanced KMs design the knowledge infrastructure that supports all specialist tracks. Coordinate with practitioners in these companion publications.

Publication	Track	Coordination Point
SL 5G	Advanced ORSA	ORSA-derived knowledge products; analytical lessons capture
SL 5H	Advanced AI Engineer	AI-assisted KM workflows; corpus design and ontology-RAG integration
SL 5M	Advanced ML Engineer	Feature data governance; training corpus quality
SL 5J	Advanced Program Manager	KM program oversight; portfolio-level knowledge governance
SL 5L	Advanced Software Engineer	Platform SDK infrastructure supporting enterprise KM systems
SL 5N	Advanced UI/UX Designer	Information architecture; knowledge portal UX
SL 5O	Advanced Platform Engineer	Knowledge platform infrastructure; search and retrieval systems

WFF Operational Consumer Note. Enterprise knowledge management exists to serve operational decision-making. 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) — are the primary users of validated lessons, TTPs, and doctrine products. The design principles in this guide — federated architecture, quality tiers, cross-domain linkage, application over capture — must all be evaluated against whether WFF practitioners can find and use the knowledge they need at the moment they need it. A KM system that serves only KMs has failed.

DISTRIBUTION RESTRICTION: Distribution authorized to U.S. Government agencies and their contractors only. Other requests must be referred to Headquarters, USAREUR-AF, C2DAO, Wiesbaden, Germany.