SPAR SOLUTIONS

Verint Engage 2026

KFAI ™ Knowledge-First AI

Participant Workbook

Why Leading CX Operations Build Knowledge Foundations First

BEFORE WE START

How to Use This Workbook

Write freely; this copy is yours to keep, and no one will collect it. Our goal today is to change how you look at and think about your organization's knowledge, and to see why making it AI-ready is both more difficult, and more valuable, than it first appears.

By the end of this session, you will be able to:

- Recognize the invisible knowledge layer and where it accumulates across an organization.
- Understand how knowledge quality determines AI success.
- Assess any knowledge article against the five attributes of AI-ready knowledge.
- Connect knowledge quality to the KPIs you already measure.
- See what an enterprise-scale knowledge diagnostic and remediation effort entails.

Today's Session at a Glance

Section	Title
1	What Normal Growth Creates
2	The Invisible Knowledge Layer
3	How Knowledge Impacts AI Initiatives
4	SPAR's Five Pillars of AI-Optimized Knowledge
5	Today's Exercise
6	Knowledge Remediation at Enterprise Scale
7	Where This Leads

1 What Normal Growth Creates

Normal growth fractures knowledge in predictable ways.

Examples of what normal growth creates

As organizations grow…	…problems emerge
Migrations and platform changes	Content scattered across systems with no single authoritative source — knowledge sprawl.
Mergers and acquisitions	Parallel knowledge bases merged without reconciliation, multiplying duplicates and conflicts.
Product and policy updates	Articles updated in parallel by different people, leaving contradictory versions in circulation — quality erosion.
Teams filling their own gaps	Shadow PDFs, team wikis, and Microsoft Teams/Slack threads stand in for the official source — governance gaps that create new silos.
Inconsistent governance	Procedures age out of date and no one is accountable for keeping them current — maintenance debt.

Most of these aren't mistakes. They're the normal byproducts of moving fast and growing. But small issues cascade: a duplicate article here, an outdated procedure there, a workaround that never makes it into the KB. Over years, they compound into system-wide costs your organization feels but can't easily see.

2 The Invisible Knowledge Layer

When a contact center's metrics look healthy, it's easy for leaders to assume they have a strong knowledge base. Average handle time is in range, first-call resolution meets expectation. The knowledge base must be in good shape.

Often it isn't. The contact center meets its metrics because people compensate and build workarounds for the information they can't find in the knowledge base. This dynamic can't be seen in the metrics themselves.

When an article is incomplete or wrong, people don't stop working. They just figure out a different way to get the job done:

- The sticky notes and cheat sheets taped to monitors.
- The answers that live only in a veteran agent's head, the person everyone asks.
- The Teams and Slack threads where the real answer got worked out last week.
- The personal docs, spreadsheets, and local files people trust more than the KB.
- The shadow SharePoint and Google Drive folders standing in for the official source.
- The escalation paths that exist, in part, to cover what the article didn't answer.

Information that lives outside the knowledge base but that an operation depends on every day is what we call the invisible knowledge layer. People know how to find it, read it, and apply it, but none of it is stored in the centralized single source of truth. It's also why strong metrics can mask a weak knowledge base.

Exercise 1 — Inventory your invisible layer

Private, just for you. Think about one team you know well. List the workarounds your people actually rely on instead of the knowledge base. Be specific: whose head, which chat, what cheat sheet.

Now the real question: if all of that disappeared tomorrow and your team could only use what's actually written in the KB, how well would they do?

3 How Knowledge Impacts AI Initiatives

AI does not fix weak knowledge. It exposes it.

For years, the problems described in the last two sections were manageable because people absorbed them. A contradiction between two articles? *"Just ask the team chat which one is right."* A procedure three versions out of date? *"I've got the current one saved on my hard drive."* A gap in the knowledge base? *"Ask the seasoned veteran."* Your people are the error-correction layer that sit between a flawed knowledge base with its gaps and the customer.

AI solutions do not have that layer. They retrieve what is in the knowledge base, and that content directly informs the responses they generate. There is no human agent in between to catch the contradiction, override the stale procedure, or supply what the article left out. Whatever is wrong in the knowledge base is what gets acted upon or shared.

The same problems we covered in the first two sections don't just remain. Their consequences become more dynamic, harder to predict, and more consequential. Sprawl raises the odds that the system retrieves the wrong article, and spends more tokens doing it. Contradictions may surface to a customer one way in one interaction and another way in the next. Outdated procedures may be acted on, creating compliance risk and liability. And the invisible knowledge layer, everything your people rely on but never put in the knowledge base, isn't available to the model at all.

> What determines whether an AI solution is trustworthy is the knowledge it can access and act on. This is why AI performance mirrors knowledge quality.

4 SPAR's Five Pillars of AI-Optimized Knowledge

Whether knowledge can be trusted and used reliably by AI depends on five key properties. At SPAR, we refer to these as the five pillars of AI-optimized knowledge. Together, they determine whether the right content can be found, interpreted, and trusted.

1. Modular
Each unit is scoped to a single intent, task, or answerable question, rather than bundling many jobs into one article.

2. Token-Structured (AI-Readable)
Each unit is formatted and written so a language model can reliably interpret its meaning, conditions, steps, and exceptions.

3. Search-Optimized
Metadata, taxonomy, and retrieval signals let the system find the right unit for the right intent.

4. Consistent
Similar content follows the same structure, terminology, and decision logic across the knowledge base.

5. Governed
Owners, standards, review cycles, and quality measures keep knowledge accurate and AI-ready over time.

The distinctions that matter:

- **Modular vs. Token-Structured:** Modular is about where one unit begins and ends. Token-Structured is about how that unit is written inside.
- **Search-Optimized vs. Token-Structured:** Search-Optimized is whether the system can find the right unit. Token-Structured is whether it can use that unit once found.
- **Consistent vs. Governed:** Consistent is whether content follows the standard. Governed is the system that enforces and maintains it.

In the next section, you'll score real articles against these five pillars.

5 Today's Exercise

The three articles below come from the same knowledge base. Score each one against the five pillars: mark every pillar Pass or Fail, and note what led you to the call.

These examples are deliberately clear-cut, so the failures are easier to see than they would be in practice. In a real knowledge base, the problems overlap, hide inside otherwise-good content, and blur together. That is what makes assessment at scale difficult.

Article A

Knowledge CenterKB-0231 • Published

Returns

Owner: (unassigned)
Updated: February 2024
Category: Returns
Tags: RTN-0231-A, x7740, SYS4-KB

SUMMARY
Covers which orders can be returned and how to process a return for a customer.

Return eligibility

Standard orders can be returned within 30 days of the delivery date. Items must be unused and in their original packaging. Final-sale items, gift cards, and personalized products are not eligible for return.

How to process a return

1. Look up the order and confirm it is within the 30-day window.
2. Verify the item condition with the customer.
3. Generate a prepaid return label from the order screen.
4. Advise the customer to ship the item within 7 days of receiving the label.
5. Once the warehouse receives and inspects the item, mark the return as complete.

Damaged or defective items

If the item arrived damaged or defective, the 30-day window does not apply. Process the return using the steps above and flag the order as “defective” so the customer is not charged a return shipping fee.

Related articles: *(none listed)*

Pillar	Pass / Fail	What led to your call
Modular	P / F	
Token-Structured	P / F	
Search-Optimized	P / F	
Consistent	P / F	
Governed	P / F	

Article B

Knowledge Center KB-0188 • Published

Stuff / Issues / Misc

Owner: (unassigned)
Updated: (no date)
Category: (uncategorized)
Tags: misc, tmp, kb-old2

SUMMARY
(none)

Use this for order and shipping problems. If a customer says their order hasn't arrived, check tracking in the order screen. If tracking shows delivered but the customer hasn't received it, have them check with neighbors and around the property, then wait 24 hours since carriers sometimes mark packages delivered early, and if it still hasn't appeared file a lost-package claim with the carrier which takes 3–5 business days, and during that time you can reship if the item is in stock, but for high-value orders over $200 get supervisor approval before reshipping. If tracking shows delayed or stuck in transit, don't file a claim yet, just have them wait until the estimated delivery date has passed by at least 2 business days. For orders that haven't shipped within the expected handling time, check if the item is backordered, and if it is, tell the customer the new ship date and offer to cancel for a full refund. If the customer wants to change the shipping address after ordering, this only works if the order hasn't entered fulfillment; once it has, it can't be changed and they have to wait or refuse the package. If they got the wrong item, apologize, generate a prepaid label to send it back, and reship the correct item free; if it's out of stock offer a refund. If a partial order arrived with items missing, check whether the missing items shipped separately first. Also returns of incorrect orders go here, see the returns process. For anything where the customer is upset or the order is over $200, consider escalating to a supervisor.

Related articles: *(none listed)*

Pillar	Pass / Fail	What led to your call
Modular	P / F	
Token-Structured	P / F	
Search-Optimized	P / F	
Consistent	P / F	
Governed	P / F	

Article C

Knowledge Center KB-0414 • Published

Refunds & Store Credit

Owner: Billing Team
Updated: April 2026
Category: Billing > Refunds
Tags: refunds, store-credit, refund-timing, eligibility

SUMMARY

Information about refunds and store credit.

Refund methods

Refunds are issued to the original payment method by default. If the order was paid with a gift card, the refund is issued as store credit. Customers may also request store credit instead of a refund to the original payment method, which is issued immediately once the return is approved.

Refund timing

Refunds to the original payment method are processed within 5–7 business days after the returned item is received and inspected. Store credit is available immediately upon approval. Credit card refunds may take an additional 2–3 business days to appear, depending on the customer's bank.

Eligibility

To be eligible for a refund, the return must be initiated within 14 days of the delivery date and the item must meet the return conditions. Customers returning items after this window may be offered store credit at the agent's discretion.

Partial refunds

If a returned item shows signs of use or is missing parts, a partial refund may be issued. Document the item condition before processing.

Related articles: *(none listed)*

Pillar	Pass / Fail	What led to your call
Modular	P / F	
Token-Structured	P / F	
Search-Optimized	P / F	
Consistent	P / F	
Governed	P / F	

What the Three Articles Were Really Testing

Every one of these articles fails on at least two pillars. Here's where, and why it matters.

	Where it fails	Why it matters
Article A	Governed, Search-Optimized	It reads cleanly, but it has no owner and was last updated over two years ago, and its bare title “Returns,” opaque tags, and absent related articles give retrieval little to match against. A solid article no one maintains or can reliably surface is still a liability.
Article B	Modular, Token-Structured, Search-Optimized, Governed	It bundles many separate problems into one unstructured wall, with a meaningless title, no metadata, and no owner. The system can neither find the right part nor use it once found. This is what an unmanaged article looks like.
Article C	Consistent, Token-Structured	The body is well-organized, but its summary tells the reader nothing, and its 14-day return window contradicts Article A's 30 days. The visible polish hides the failures that matter.

Two of these articles fail on Token-Structured, but for completely different reasons: Article B is an unstructured wall, while Article C is a clean article with an empty summary. The same pillar, two different failures.

And catching Article C's return-window error meant remembering what Article A said. That contradiction is invisible if you read either article alone. It only appears when you compare them, which is exactly why assessing a knowledge base gets harder as it grows.

6 Knowledge Remediation at Enterprise Scale

You just scored three articles, and they were built to be clear-cut. Most organizations hold anywhere from hundreds to tens of thousands of knowledge articles, and in a real knowledge base the failures are rarely this easy to spot. Assessing knowledge at that volume is a major undertaking, and it demands a level of rigor and detail that spot-checking a few articles cannot provide.

Start with the assessment itself. Every article has to be evaluated against all five pillars, one by one. But the most damaging problems do not live inside individual articles. They live between them: the contradiction you caught between Article A and Article C only exists when the two are compared. Finding those requires checking articles against each other, not just reading them in isolation, and the number of possible comparisons grows far faster than the number of articles.

Then comes remediation, which is the larger and harder portion of the work. You cannot remediate without diagnosing first, because you cannot fix an article correctly without knowing what is wrong with it, what it overlaps, and what depends on it. And fixing one article is rarely a contained change. Splitting an overloaded article into properly scoped units means deciding where each piece belongs. Resolving a single contradiction means finding every other article that touches the same topic and reconciling them to one source of truth. Each change ripples outward, because a knowledge base is an interconnected system, not a stack of independent documents.

That is what makes this work compound. A few hundred articles, each assessed against five pillars and against one another, each remediated in a way that affects its neighbors, is not a task a small team can clear in a sprint. It is a sustained, systematic effort, and it is the reason knowledge does not get fixed by simply asking people to tidy it up in their spare time.

7 From Diagnosis to Governance: A Repeatable Approach

The scale and complexity of this work are exactly what Knowledge-First AI, or KFAI™, was built to handle. It is SPAR's structured approach to making an organization's knowledge AI-ready, and keeping it that way. It works in three connected stages: diagnosis, remediation, and governance.

Diagnosis

The diagnosis measures the current state of the knowledge base against the five pillars. Each pillar has its own set of diagnostic tools, so the assessment is systematic and repeatable rather than a matter of judgment, and it covers the full knowledge base rather than a sample of it.

Remediation

Remediation fixes what the diagnosis finds. Every diagnostic tool is paired with a corresponding remediation tool, so each problem that is identified has a defined path to being resolved. Diagnosis tells you what is wrong and why; remediation does the work of making it right.

Governance

Governance is what keeps the knowledge base AI-ready after the first pass. Knowledge decays as policies change, and new knowledge is created all the time. Because diagnosis and remediation are paired and tool-driven, the same machinery can be run on a regular cadence. Governance is not a separate program; it is the ongoing operation of the same tools that found and fixed the problems in the first place.

One more thing worth saying plainly. The work that makes knowledge “AI-ready” is the same work that makes it better for the people who use it today. Cleaner, single-purpose, consistent, well-governed content is easier for a human agent to find and trust, and easier for a customer to be served from. The improvement shows up in the metrics you already track, before a single AI initiative is switched on. AI-readiness is the reason to start; the human-side improvement is the return you see along the way.

Where this leads

If you would like to know where your own knowledge stands, talk with us while you are here. We can walk you through what a diagnosis looks like and what it tends to surface, and help you judge whether this is work worth doing for your organization.