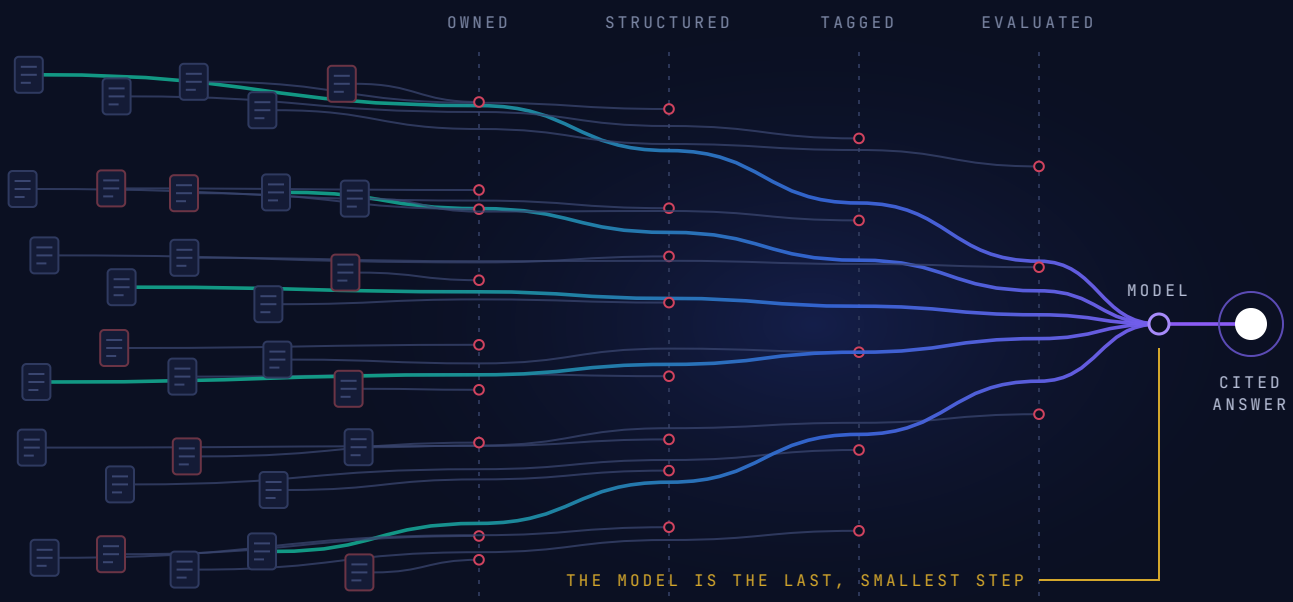
f/cx.ai

WHITEPAPER · 2026

KNOWLEDGE ARCHITECTURE · CONTACT CENTER · CONTENT DESIGN

Knowledge Is the Product

Why enterprise AI performance is usually constrained by knowledge architecture, not by the model.



PDFS · WIKIS · MACROS · POLICY DECKS · EMAIL

Skip Vanderburg
Founder & Chief Experience Officer

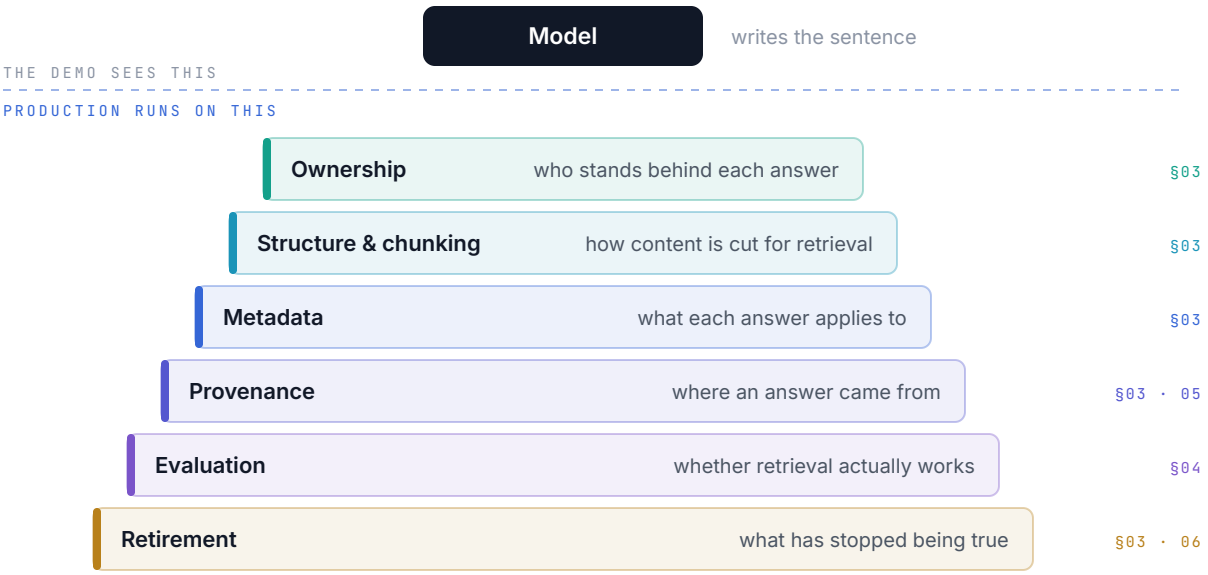
FRACTIONALCX.AI
INNOVATE@FRACTIONALCX.AI

THE ARGUMENT

The model is rarely the bottleneck

When an AI agent quotes a retired fee or invents a policy, the review rarely finds a weak model. It finds an article nobody owned, a PDF cut in half and two versions that disagree. The model answered faithfully. The knowledge was wrong.

FIGURE 1 · WHAT ACTUALLY DECIDES THE ANSWER



The budget goes to the top layer. The answer is set by the six below it. Swap the model and every layer underneath still decides what the customer hears.

Who this is for

CIOs and CTOs funding AI in service, contact-center leaders who own the outcome, and the knowledge and service-operations teams who own the content underneath.

What it is not

Not a RAG tutorial. It is the operating discipline that makes any retrieval stack trustworthy: ownership, structure, evaluation and the cost case to fund them.

61%
of service leaders have a backlog of knowledge articles awaiting edits
GARTNER, 2024

63%
of organizations lack, or are unsure of, data practices ready for AI
GARTNER, 2025

42%
say over half their AI projects stalled or failed on data readiness
FIVETRAN, 2025

4×
AI customer service fails at nearly four times the rate of other AI uses
QUALTRICS XM, 2025

Written for people, read by machines

A person reads an article top to bottom, skips what doesn't apply and asks a colleague when it's unclear. A retriever reads fragments, out of order, with no colleague.

FIGURE 2 · ONE ARTICLE, TWO READERS

AS A PERSON READS IT

Refunds & Returns

See above for exceptions.

* Excludes Gold members. See 2023 memo.

→ CUT INTO 500-TOKEN CHUNKS

AS THE RETRIEVER GETS IT

CHUNK 3
Table cut mid-row. The fee is in one chunk, the plan it applies to in another.

CHUNK 4
"See above for exceptions." In a fragment, nothing is above.

CHUNK 6
The exception lives in a 2023 memo that was never indexed.

CHUNK 9
The 2024 version is indexed too. Nothing says which is current.

Nothing here is badly written. It was written for a reader who has the whole page. Illustrative.

FIGURE 3 · SIX WAYS GOOD PROSE BREAKS RETRIEVAL

01
Buried answers
The answer sits in paragraph six, under a heading about something else.

02
Context by reference
"As above," "this plan" and "see policy" mean nothing in a fragment.

03
Tables and images
Fees in a screenshot, or a table split across chunks.

04
Conditional truth
One article, five customer types, exceptions in footnotes.

05
Duplicates and drift
Three versions, all indexed, none marked current.

06
No owner, no date
Nothing says who stands behind it, or until when.

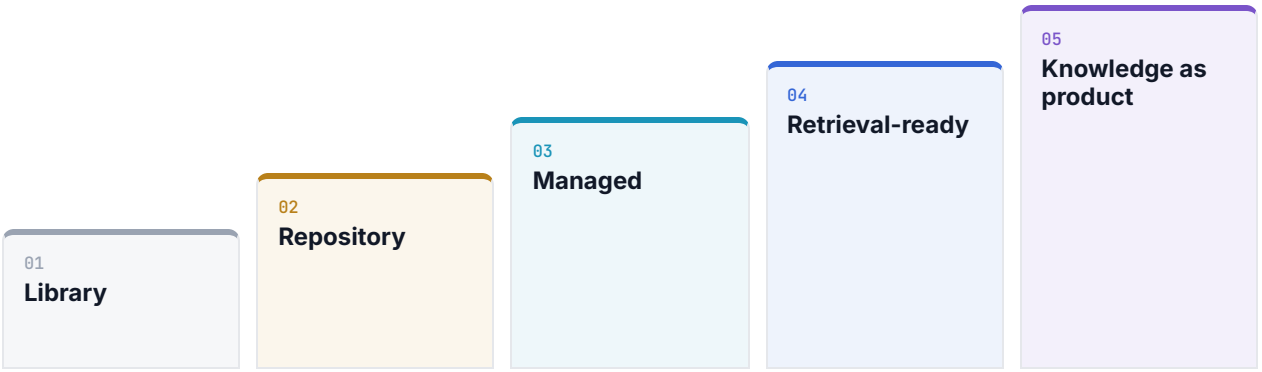
THE LAW HAS ALREADY DECIDED

In **Moffatt v. Air Canada** (2024), the chatbot said a bereavement fare could be claimed after travel. The airline's own policy page said it could not. The tribunal held the airline to the chatbot's answer. **Two sources disagreed, and the customer was owed the one he was told.**

Five levels of knowledge readiness

Readiness is visible from outside: who owns an answer, how it is cut, and whether anyone can say which version answered. Most contact-center estates sit at level two.

FIGURE 4 · FIVE LEVELS OF KNOWLEDGE READINESS



01	02	03	04	05
Written for people, found by search.	Centralized. Owners unclear.	Owners, templates, review dates.	Answer units, metadata, golden set.	Backlog, owner, cost, tests in the pipeline.

MOST CONTACT CENTERS ARE HERE · 02

THE STEP THAT MATTERS · 03 → 04

FIGURE 5 · SIX DIMENSIONS, BY LEVEL

	01	02	03	04	05
Ownership	Author left	Team-owned	Named owner	Owner per intent	Owner with SLA
Structure	Long pages	Templates	Headed sections	Answer units	Modeled objects
Metadata	None	Category tags	Product, audience	Intent, scope, dates	Governed schema
Provenance	Unknown	Last edited	Version history	Source on every chunk	Cited in every answer
Evaluation	Complaints	Search clicks	Article ratings	Golden set by journey	Gates every change
Retirement	Never	Ad hoc	Annual review	Review-by dates	Auto-expire, alerts

Your level is your lowest row, not your average. A golden set cannot rescue content nobody owns, and an owner cannot fix what nobody measures.

> **The model is a commodity you rent. The knowledge is the asset you own, and the only one a competitor cannot copy.**

Treat every answer as a managed object

Retrieval-ready knowledge is not a better wiki. It is a set of small, owned objects, each carrying what a machine needs to use it safely.

FIGURE 6 · ANATOMY OF A KNOWLEDGE OBJECT

ANSWER UNIT · KB-2291

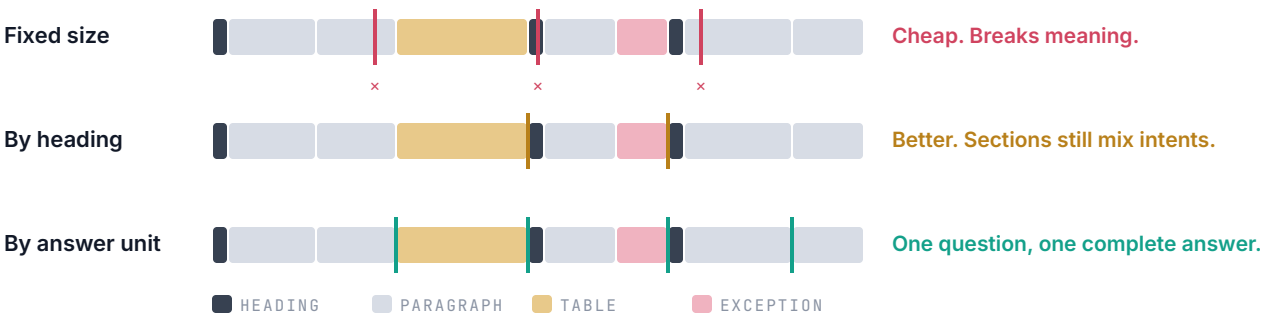
Refund window for annual plans

"Annual plans can be refunded in full within 30 days of purchase or renewal."

OWNER	Owner Backup	Billing Policy Lead Service Knowledge Team	One named owner per answer, not per article.
CHUNK	Unit Length	One intent, self-contained One complete answer	Cut by question, not by length.
METADATA	Intent Scope	refund.annual_plan Annual plans · US · all channels	Tags describe who it is true for.
PROVENANCE	Source Approved	Refund Policy v3.2 §2 Aug 14, 2026 · Legal	Every answer can name its source.
RETIREMENT	Review by Expires	Feb 14, 2027 When policy v4 publishes	No date, no publish.

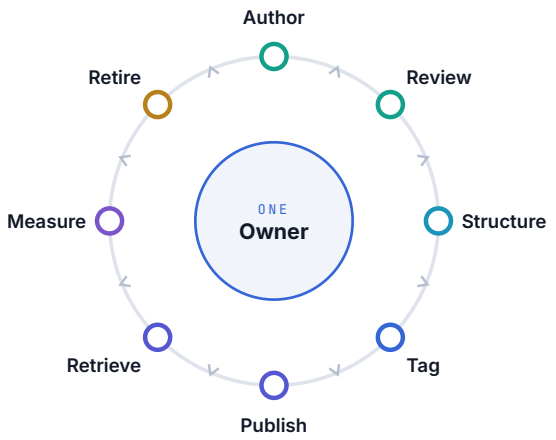
Five disciplines, one object. If a field cannot be filled, the answer is not ready for an agent. Illustrative.

FIGURE 7 · THREE WAYS TO CUT THE SAME ARTICLE



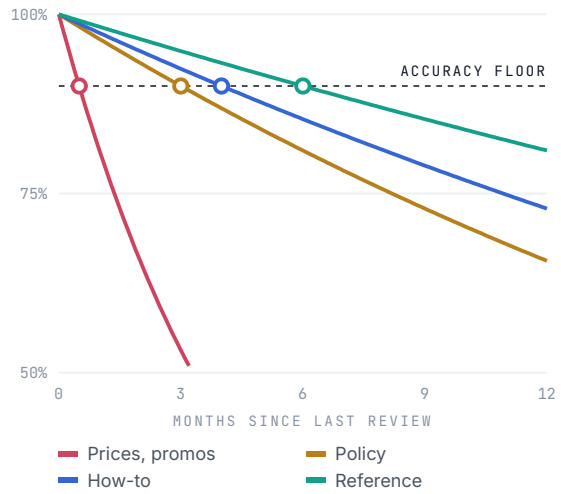
Chunking is a content decision, not a parameter. The best cut follows the questions customers ask.

FIGURE 8 · THE KNOWLEDGE LIFECYCLE



Retire is a stage, not an event. The loop only closes when measurement can send an answer back to its owner.

FIGURE 9 · HOW FAST KNOWLEDGE DECAYS



Set review by volatility, not by calendar. Each dot is when that tier drops below the floor. Illustrative.

REVIEW TRIGGER BY TIER

Prices & promotions
Expires with the offer

Policy & eligibility
Every policy change, plus quarterly

Procedures & how-to
Every product release

Reference facts
Twice a year

Three roles most estates do not have yet

IN SERVICE OPERATIONS

Knowledge Product Owner

Owns the estate like a product: backlog, owners, review cadence and the cost of getting it wrong.

IN CONTENT DESIGN

Answer Designer

Rewrites articles into answer units, sets the metadata schema and keeps conditional truth explicit.

IN TECHNOLOGY

Retrieval Engineer

Owns chunking, indexing, the golden set harness and the thresholds that decide when to answer.

WORTH NOTICING

Over a third of service leaders have **no formal process for updating outdated articles** (Gartner, 2024). By 2026, **58%** planned to upskill agents into knowledge management specialists. The role is arriving. The ownership model usually is not.

Test retrieval the way customers ask

Benchmarks test whether a system can find facts. Customers test whether it finds *their* fact, mid-journey, in their words. Build the test set from journeys, not from the article list.

FIGURE 10 · FROM JOURNEY TO GOLDEN SET

JOURNEY · MOVING HOME

01 Plan the move	02 Transfer service	03 Check coverage	04 Final bill	05 Dispute a charge
"Can I take my plan with me?"	"Will there be a gap?"	"Do you cover my new zip?"	"Why am I billed twice?"	"I got a moving fee."
"How early do I tell you?"	"Do I need a technician?"	"Same speed there?"	"When does the old one stop?"	"That's not what I was told."

↓ EVERY QUESTION BECOMES A TEST CASE

QUESTION, IN CUSTOMER WORDS	CUSTOMER CONTEXT	EXPECTED SOURCE	EXPECTED ANSWER
"Can I take my plan with me?"	Annual plan · moving in-region	Moving Home Policy v2 §1	Yes. Terms carry over in-region.

Questions come from transcripts, not from authors. Forty journeys at twenty-five questions each is a thousand-case golden set. Illustrative.

FIGURE 11 · FIVE GATES EVERY ANSWER MUST PASS

	SHARE OF TEST CASES PASSING	WHERE THE FIX LIVES
Retrieved Right source in the top five	82%	Coverage, metadata
Ranked Right chunk ranked first	68%	Chunking
Grounded Answer supported by the chunk	61%	Thresholds, prompt
Cited Citation points to that source	58%	Provenance
Right for them True for this customer's context	49%	Scope metadata

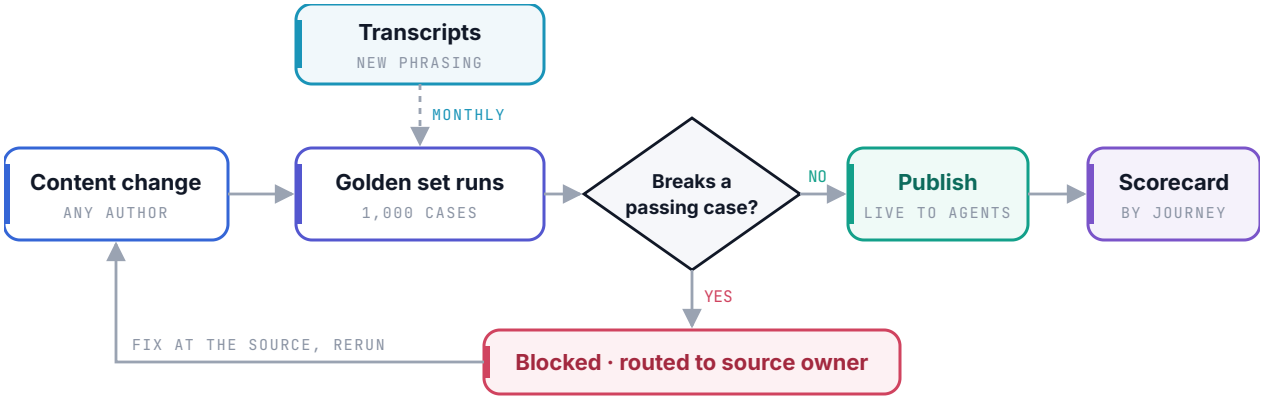
Each gate fails for a different reason, and has a different owner. A single "accuracy" number hides which one. Illustrative.

FIGURE 12 · THE SCORECARD, BY JOURNEY

	RETRIEVED	RANKED	GROUNDED	CITED	RIGHT FOR THEM	WHY
Billing dispute	88%	74%	70%	66%	52%	Exceptions in footnotes
Plan change	91%	83%	80%	78%	71%	Healthy
Moving home	84%	70%	66%	63%	38%	Missing region scope
Service outage	62%	58%	52%	50%	46%	Stale status pages
Cancel	90%	79%	74%	71%	60%	Retention offers drift
Device setup	86%	54%	51%	49%	47%	Long manuals, cut by size

The average hides it. Three weak journeys, three different fixes. Illustrative.

FIGURE 13 · THE GOLDEN SET AS A RELEASE GATE



Run it like a regression suite. Failures route to the owner of the source, not to the AI team.

17–34%

EVEN SPECIALIST TOOLS MISS

Leading retrieval-based legal AI tools gave incorrect information this often in Stanford testing. **Retrieval is not solved by buying it.**

STANFORD REGLAB, 2024

FIGURE 14 · FOUR NUMBERS TO PUBLISH EVERY WEEK

Retrieval recall
Right source in the top five results
0 BAND 80–90% 100

Groundedness
Answers fully supported by retrieved text
0 BAND ≥ 95% 100

Citation accuracy
Citations that point to the true source
0 BAND ≥ 98% 100

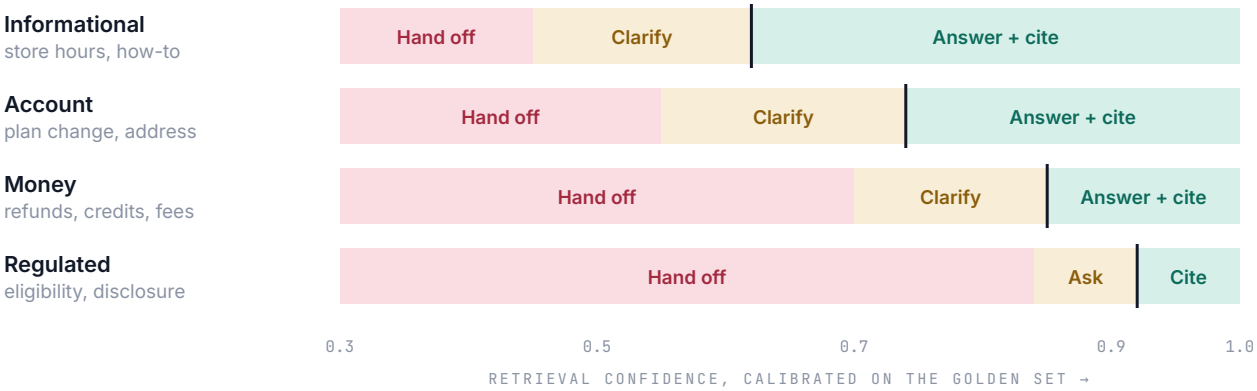
Resolution that holds
No repeat contact within seven days
0 BAND 68–75% 100

Bands, not targets. A band invites the question “why did it move?” A target invites gaming. Illustrative bands.

Answer, ask or hand off

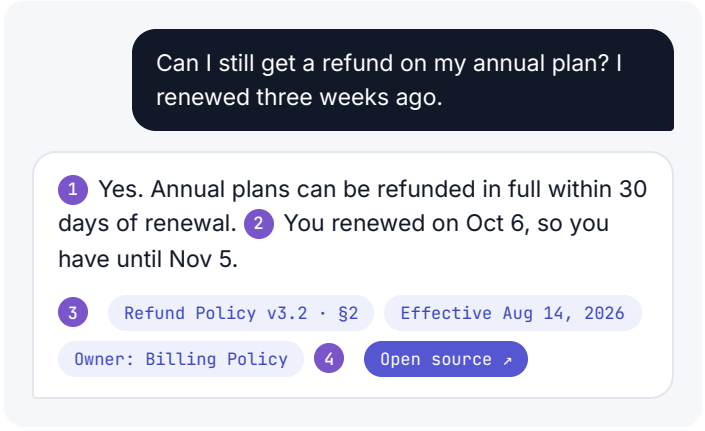
A retrieval-ready estate knows when it does not know. Thresholds decide whether the agent answers, asks a clarifying question or hands off, and the bar rises with the consequence.

FIGURE 15 · THREE ZONES, SET BY CONSEQUENCE



The answer line moves right as the stakes rise. A regulated answer needs near-certainty and a verbatim source. Illustrative.

FIGURE 16 · ANATOMY OF A CITED ANSWER



1. **Answer first**
   The direct answer leads. Detail follows.
2. **Context applied**
   Customer data narrows a general rule to their case.
3. **Source, version, date**
   Provenance travels with the answer.
4. **Openable**
   The customer and the agent can read the source.

No citation, no answer

For money and regulated tiers.

Calibrated, not chosen

Set on the golden set. Rechecked after every re-index.

Handoffs carry sources

The person starts where the AI stopped.

Clarify by design

One good question beats a confident wrong answer.

Put a price on a wrong answer

Bad knowledge is expensive in places no knowledge budget sees: repeat contacts, escalations, credits, churn and agents searching. Price it once and the remediation case writes itself.

FIGURE 17 · ONE YEAR OF BAD KNOWLEDGE, PRICED



INPUTS · 1.2M AI CONVERSATIONS · 5% FAIL ON KNOWLEDGE · 250 AGENTS SEARCHING 15 MIN A DAY

The two largest costs never appear on a knowledge budget. Churn lands in marketing. Search time lands in operations. Illustrative.

FIGURE 18 · THE FORMULA

Knowledge failure rate
from golden set and transcripts

×

Volume by journey
AI and assisted contacts

×

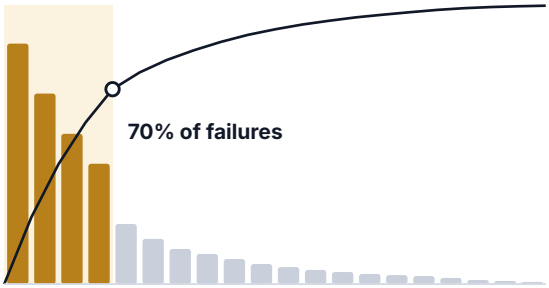
Cost per failure
by consequence tier

+

Search time
agent minutes × loaded rate

Price by journey, not in total. The fix list falls out of the same table.

FIGURE 19 · WHERE THE COST CONCENTRATES



EACH BAR = 5% OF ARTICLES, RANKED BY FAILURES CAUSED

Fix the top fifth first. In most audits a small share of articles drives most failures. Illustrative.

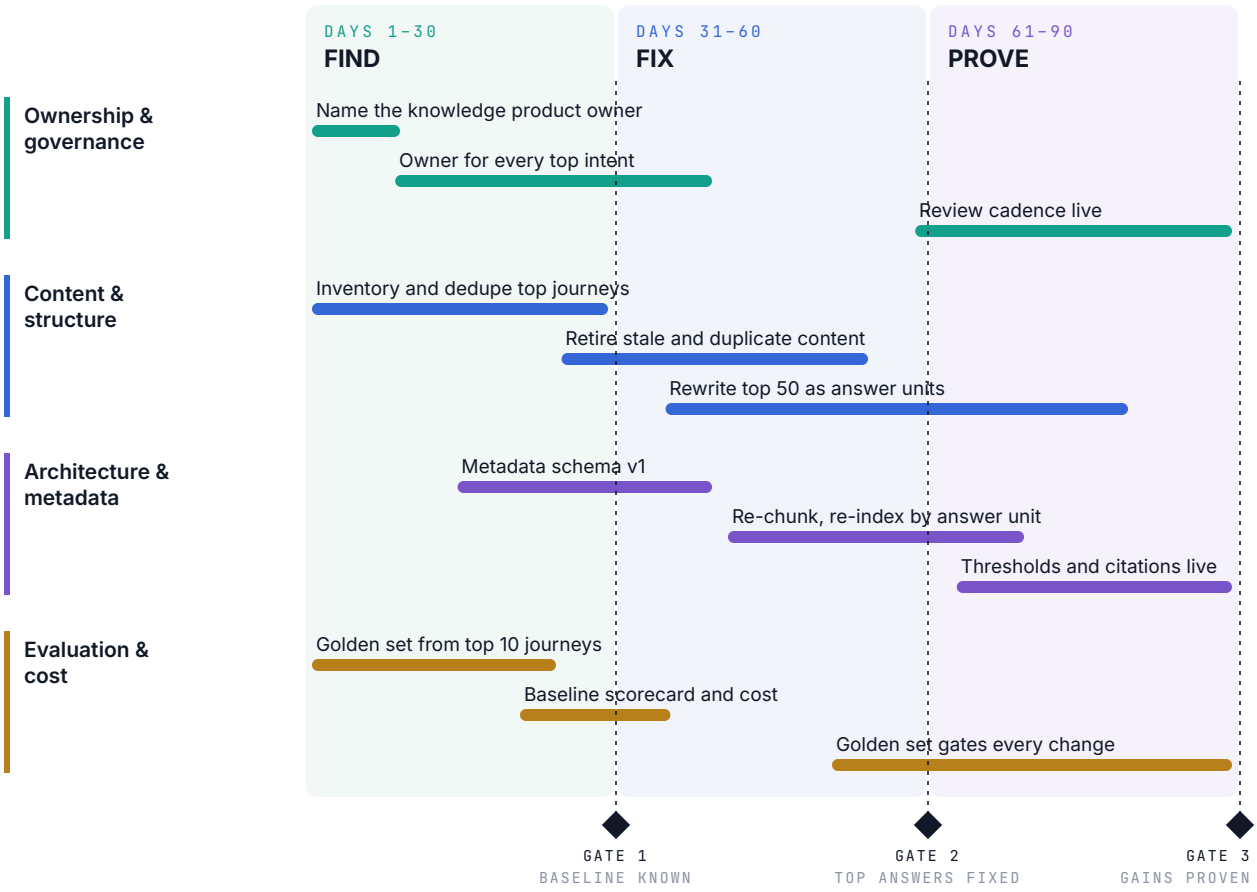
THE CLOCK IS RUNNING

Gartner predicts organizations will abandon **60% of AI projects** unsupported by AI-ready data through 2026. The cost of bad knowledge is not only what it costs today. It is the program it cancels.

Ninety days, three gates

Remediation is not a migration. Start with the journeys that cost the most, fix the answers they depend on, and prove it on the golden set before widening.

FIGURE 20 · A 90-DAY KNOWLEDGE REMEDIATION ROADMAP



Measure first, then fix. Teams that rewrite before they baseline cannot show what changed. A starting plan to adapt.

FIGURE 21 · WHAT YOU HAVE ON DAY 90

- **A named owner** for every top intent
- **Top 50 answers** rebuilt as answer units
- **A metadata schema** enforced at publish
- **A golden set** gating every change
- **Thresholds and citations** live by consequence tier
- **The cost of bad knowledge** measured before and after

Where to start

Three seats share one answer. Put them in a room with ten questions and you will know which section of this paper to start with.

FIGURE 22 · THREE SEATS, ONE ANSWER

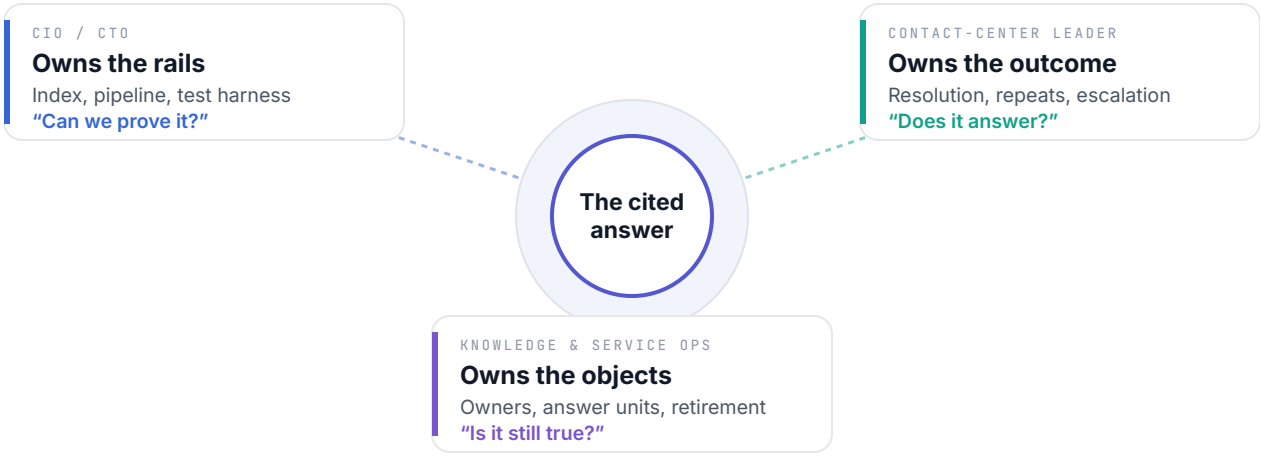


FIGURE 23 · TEN QUESTIONS TO ASK ON MONDAY

01 Top 50 questions by volume? §04	02 Owner for each answer? §03	03 How many refund policies indexed? §01	04 What happens to a table? §03	05 Which source caused Tuesday’s miss? §03
06 What happens at low confidence? §05	07 Do money answers always cite? §05	08 When did you last retire content? §03	09 Does a change run a test first? §04	10 What did bad knowledge cost? §06

0–6 CLEAR ANSWERS · THE NORMAL STARTING POINT

7–10 · READY TO WIDEN

Tick the ones you can answer, with a name. Each blank points to the section of this paper to start with.

> **The enterprises that win with AI in service will not have the best model. They will have the best-kept answers.**

f/cx.ai

WHERE THIS GOES NEXT

Fix the knowledge. Then let the agent answer.

f/cx.ai is a customer experience design and innovation agency for the AI era. We design, prototype and prove the journeys, service models and AI agents enterprise organizations run on.

CONTACT CENTER MODERNIZATION

Rebuild the estate

Knowledge, routing and AI in the service path, redesigned around the journeys that cost the most.

AGENT ASSIST

Answer the agent first

Cited answers inside the agent desktop, so people and AI draw on one source of truth.

IA & CONTENT DESIGN

Make it retrieval-ready

Answer units, metadata schemas and golden sets, designed to be owned and maintained.

The model writes the sentence.
The knowledge decides whether it is true.

THE CONVERSATION WE HAVE MOST OFTEN

Why the agent was right in the demo and wrong in production, and what has to be true of your knowledge before you widen what it may answer.

CONTACT

Skip Vanderburg
Founder & Chief Experience Officer

innovate@fractionalcx.ai
fractionalcx.ai

© 2026 FRACTIONALCX.AI LLC · ALL FIGURES ORIGINAL · SAMPLE DATA IN FIGURES IS ILLUSTRATIVE AND NOT ANY NAMED CLIENT'S RESULTS

Where every external number came from

Every figure in this paper is original. External data points are attributed where they appear and listed here in full. Vendor-commissioned studies are noted.

KNOWLEDGE AND SERVICE OPERATIONS

Gartner 2024	Customer service leaders survey, press release, Dec. 9, 2024. 61% report a backlog of knowledge articles awaiting edits; over a third lack a formal process for updating outdated articles; 85% plan to explore or pilot customer-facing conversational GenAI in 2025. n = 187, fielded July–Aug. 2024.
Gartner 2026	Customer service leaders survey, press release, Feb. 18, 2026. 91% report executive pressure to implement AI; 58% aim to upskill agents into knowledge management specialists. n = 321, fielded Oct. 2025.

DATA READINESS AND AI OUTCOMES

Gartner 2025	"Lack of AI-Ready Data Puts AI Projects at Risk," press release, Feb. 26, 2025. 63% of organizations do not have, or are unsure if they have, the right data management practices for AI (n = 1,203 data management leaders, July 2024). Prediction: through 2026, organizations will abandon 60% of AI projects unsupported by AI-ready data.
Fivetran 2025	AI and Data Readiness Survey, conducted by Redpoint Content, Q1 2025. 42% say more than half of their AI projects were delayed, underperformed or failed due to data readiness. n = 401 data leaders. Vendor-commissioned.
Qualtrics XM Institute 2025	2026 Consumer Experience Trends research. AI-powered customer service fails at nearly four times the rate of other AI uses. 20,000+ consumers in 14 countries, fielded Q3 2025.

RETRIEVAL ACCURACY

Stanford RegLab & HAI 2024	Magesh et al., "Hallucination-Free? Assessing the Reliability of Leading AI Legal Research Tools." Lexis+ AI produced incorrect information more than 17% of the time; Westlaw AI-Assisted Research more than 34%. Retrieval failures cited among causes.

LAW AND LIABILITY

Moffatt v. Air Canada 2024	2024 BCCRT 149, British Columbia Civil Resolution Tribunal. The airline was held responsible for its chatbot's bereavement-fare answer, which contradicted its own policy page.

Prepared by f/cx.ai, a customer experience design and innovation agency for the AI era. Sample data inside figures is illustrative and is not any named client's results. The maturity model, gates, thresholds and cost model are working instruments to revise against your own evidence. This paper is not legal advice.