Straive

EXCLUSIVE CHECKLIST REPORT

Trusted Knowledge Meets GenAI in CX

KMWorld
KMWorld.com

Customer Trust Depends on Well-Governed, Curated Knowledge

By Marydee Ojala, Editor-in-Chief, *KMWorld*

Customer service proves to be a perfect area to demonstrate the value of newer AI technologies. As expectations from customers ratchet up, customer service is a flashpoint for designing interfaces that meet customer needs and enhance the customer experience (CX). Trust is integral to interactions with customers, particularly as opinions about trustworthy AI differ—which provides the ideal opportunity for knowledge management to shine.

Speed is certainly one aspect of today's customer expectations. Wait a day for an answer to your question? Not in the cards. It's getting to the point that even a minute is too long for attention-span-challenged customers. But getting an almost instantaneous response isn't the only requirement. Customers want a system that remembers them, is courteous and polite, and returns an accurate, reliable answer. Is that too much to expect? Is magic involved?

Sometimes it seems that customers *do* expect magic. They want instant answers and consistency regardless of whether they're on chat, email, or the phone. What's worse, they are comparing this experience with all the others they've ever had. Too many times, companies fall short, lacking the requisite magic customers expect. With Generative AI being touted as a solution to, well, just about everything, the expectations are higher than ever. GenAI, however, isn't magic, although it often gives a good imitation of it.

FRAGMENTED DATA

The problem lies in the fragmentation of data inside organizations. The perennial problem of data residing in silos still exists. Needed information, such as policies, product details, and who's entitled to what, is scattered across dozens of systems. Human agents scramble to find which silo has the data the customer is waiting for. And while that happens, the customer continues to wait, probably unhappily. To make matters worse, the customer service agent may decide to hand the request over to someone else, possibly in another area of the organization or to a more senior agent. Too often, that means the customer has to repeat everything to the new person. Starting over does not align with customer expectations, extending the time to resolution and frustrating the customer. Trust evaporates.

Lately the notion of GenAI as magician has become popular, even if it's not phrased as being magical. Just throw AI at customer service, deploy some chatbots, launch some AI assistants, and all will be well. But it takes more than that. It takes governance. It takes curation. It takes involvement from subject matter experts. It takes fixing underlying knowledge bases that contribute to fragmented data. It takes patience. It takes human intervention. With well-governed, curated knowledge, the CX with an organization can look magical.

If foundational knowledge is a mess, throwing GenAI at it without thinking through the ramifications may worsen the situation, not improve it. In too many organizations, duplicate articles contradict each other. Ownership either doesn't exist or is unclear; edits happen without a paper trail, not even a digital paper trail. New hires struggle to figure out which piece of information is the correct one. Even those who aren't new hires may have problems sorting that out. Multiple channels exist but provide differing answers to queries, leading not to clarification but to confusion. The result is detrimental to both CX and the employee experience.

Given such a chaotic environment, introducing AI into the mix exacerbates the problems. Automating confusion should not be the object of the exercise. Customer service agents should not have to wonder if what is in a knowledge base is actually true. Consistency should be a given. When metadata is missing or inconsistent, search becomes a guessing game that GenAI amplifies rather than fixes. What about audit trails? Are they being implemented in a robust fashion or are they more likely to not even survive a gentle breeze? At what point does knowledge cease to be an asset and become a liability?

FIXING KNOWLEDGE

To ensure that knowledge is, indeed, an asset for an organization, knowledge managers need to embrace the notion of governance. This might not be the stuff of magic and it might not even happen instantaneously, but it is a necessary piece of building trust. Governance that engenders trust in the data relies on basic principles. It requires determining ownership, editorial standards, metadata tags, and release cycles. Once that is accomplished, adding GenAI brings benefits such as detecting customer frustration and replying to queries in a human-like fashion, not to mention providing speedy responses.

In today's world, customers expect their personal data to be private. They hear too much in the news about data breaches and want to be reassured that, when they contact a customer service agent, it's a responsibility of the company to protect the privacy of the customer. Privacy filters are essential to ensuring that GenAI operates within the boundaries of trust and compliance. It's important that they automatically identify and redact any personally identifiable information or sensitive data before processing a request.

For example, Straive's Senior Vice President, Priyanka Raju told me, "if a customer inquiry contains an account number or medical reference, the system replaces that content with a secure placeholder (such as '[ACCOUNT_NUMBER]') prior to generation. The AI then uses governed knowledge sources to craft an appropriate response, and the actual data is reintroduced only during secure post-processing. This approach allows organizations to scale AI-assisted service safely, maintaining both efficiency and regulatory integrity."

STANDARDIZING KNOWLEDGE ACROSS REGIONS

When considering how to standardize across geographic regions, particularly in light of possible local policy nuances, balancing standardization with local requirements can be tricky. Raju sees the key to this as "standardizing the framework, not the flavor." To use a meal metaphor, lunch might be the framework, but the flavor could be sandwich, taco, soup, curry, sushi, or bratwurst.

"An understanding of context enables agentic AI to ascertain which piece of information from a knowledge base is the most relevant for a particular query. It also allows for a more proactive approach to customer service and an improvement in CX overall. Agentic AI encourages a transition from a passive experience where customer service is merely responding upon request to a more active experience where customer needs are anticipated."

A customer who orders the bratwurst but is served sushi is not going to be pleased. In knowledge management terms, rather than food ones, think about establishing uniform templates, taxonomies, and editorial standards globally, and then adapting them through metadata-driven localization.

Metadata elements could be region, product, policy version, and effective dates. This allows content to adjust dynamically to customer context. Governance bodies, such as centralized service centers or global capability centers, oversee this balance. At Straive, for example, said Raju, such governance frameworks ensure that while editorial and structural standards remain universal, local policy and linguistic nuances are fully respected. This model enables a seamless, globally consistent experience that is rooted in regional authenticity.

INFLECTION POINTS

Change is coming rapidly to customer service interactions. Straive points to changes in customer expectations as a major inflection point. But there are others. The adoption of GenAI within enterprises, particularly with advances in language models, machine learning, predictive analytics, and search technologies, is a potent inflection point. Agentic AI contributes to this as well. AI-powered agents, with their ability to autonomously perform tasks that enhance the performance of human agents, add a new dimension to CX that should delight customers accustomed to a more pedestrian experience. While GenAI's forte is interpreting what customers want, agentic AI actually does things, such as check entitlements, update orders, and schedule callbacks, within established policy guardrails.

Agentic AI moves CX to a higher level, particular when it includes AI-powered self-service and context-aware responses. Customers can receive immediate, accurate, and conversational answers 24/7; no recorded message saying the department is closed and will reopen at 9 a.m. An understanding of context enables agentic AI to ascertain which piece of information from a knowledge base is the most relevant for a particular query. It also allows for a more proactive approach to customer service and an improvement in CX overall. Agentic AI encourages a transition from a passive experience where customer service is merely responding upon request to a more active experience where customer needs are anticipated.

Fraud detection and prevention provide some excellent examples of proactive initiatives. Not that long ago, I was notified of a suspicious charge on my credit card. The card issuer's system identified the charge as potential fraud almost immediately, certainly before I was aware of it. I confirmed it wasn't authorized by me and the charge was instantly removed. All this without ever speaking to a human being. Yes, happy customer here.

Similar types of proactive initiatives include payment failures, delayed deliveries, and login difficulties. This signifies a shift from customer service solving a problem when asked to do so towards a different approach, of noticing a problem and solving it before being asked to do so. Again, it looks like magic, but to make it work, knowledge must be structured so that it can be acted upon. Clear triggers, audit trails, and version control need to exist before invoking AI.

Unified knowledge platforms are another inflection point that increases customer satisfaction. When knowledge is siloed, it's difficult to ensure consistency. Plus, customers have a variety of devices available to them, which they use to contact customer service. Omnichannel support should be seamless across the unified knowledge platform.

THE PROMISE AND PERILS OF GENAI

GenAI promises much. In tandem with a robust governance system and curated knowledge stored in non-siloed knowledge hubs, GenAI delivers on that promise. However, it's a prevalent misconception to think that GenAI can independently transform customer service and bring CX up to the level expected by today's customers. Believe this at your peril. Too many organizations deploy AI tools expecting immediate gains, without first establishing a robust foundation of governed, accurate knowledge. When the underlying content is fragmented or outdated, AI only amplifies the inconsistencies.

Straive has produced an excellent 8-step roadmap for organizations looking for a path towards the trusted knowledge that is essential to creating an excellent CX and positively feeding GenAI and agentic AI initiatives. The roadmap outlines a disciplined sequence for organizations to follow.

In practice, the success of GenAI depends entirely on the quality and governance of the knowledge it draws upon. Organizations that achieve sustainable results treat knowledge as a managed product—with defined ownership, editorial standards, metadata, and release cycles. Once this foundation is in place, GenAI becomes a force multiplier: accelerating retrieval, generating accurate and brand-aligned responses, and guiding consistent next steps across all channels. Not only that, it becomes scalable as CX grows.

Keep in mind that your CX is only as good as the knowledge feeding it. GenAI and agentic AI speed the process along but a strong underlying KM system, based on robust governance and curation, is essential. There really aren't any shortcuts to this and no magic is involved. The bottom line is: Fix the foundation first. Then watch how CX improves and marvel at the impact gained. ▮

Marydee Ojala is editor-in-chief of *KMWorld* magazine and is conference program director for Information Today, Inc. She works on conferences such as Enterprise Search & Discovery, KMWorld, Data Summit, and AI & ML Summit. She is a frequent speaker at U.S. and international information professional events. In addition, she moderates the popular KMWorld webinar series.

Trusted Knowledge Meets GenAI in CX

By Priyanka Raju, Senior Vice President, Straive

Introduction: The CX Inflection Point

Customer expectations have shifted for good. People now expect quick answers, a respectful tone, accurate information, consistent experiences across channels, and recognition of their prior history. Each interaction is judged against the best service they have ever received. Many service organizations struggle to meet this bar because the information needed to answer questions—policies, entitlements, product details, and regional rules—sits in disconnected systems. Agents spend time searching instead of solving. Customers repeat themselves during handoffs. Resolution takes longer, and trust declines.

A practical path forward is to treat the knowledge base as the primary service asset and apply generative AI (GenAI) that is explicitly tied to governed content. Subject-matter experts keep information current. GenAI speeds retrieval, drafts responses, and guides next steps. Operations teams measure performance and feed improvements back into content and workflows.

The result is a fast, reliable service model without added risk. The organization manages knowledge like a product, grounds AI in that knowledge, and links outcomes to the content and processes that produced them. When agents can trust what they see—and have assistive tools that adapt to context—customers feel heard and receive clear next steps. When leaders can trace outcomes to sources, they can strengthen governance where it matters.

THE KNOWLEDGE CHALLENGE IN CX

Unmanaged knowledge causes many service failures. Multiple repositories create duplicates and contradictions. Edits happen ad hoc, with no audit trail. New hires face long ramps because no single source can be trusted. Customers receive conflicting answers depending on the channel or representative. Core systems—CRM, entitlement, and knowledge—are not connected at the moment of service, so agents must piece together context while the customer waits.

Key gaps often include:

- ✓ **Ownership and cadence:** No clear owners or review cycles; updates drift out of sync with product and policy changes.
- ✓ **Editorial standards:** Tone and structure vary; references to source policies are inconsistent or missing, reducing trust and metadata richness.
- ✓ **Metadata:** Articles lack tags for region, product, entitlement, policy version, and effective dates, leading to imprecise retrieval and risky automation.
- ✓ **Change control:** Audit trails are incomplete; refresh dates aren't visible; content releases lag product or policy changes.
- ✓ **Feedback loop:** Frontline input is fragmented or lost; fixes reach production slowly; contributors rarely see what changed, so issues recur.

When outdated or contradictory knowledge feeds bots or GenAI models, errors scale up quickly. The core challenge is scale and fragmentation. Many organizations do not manage knowledge with the same rigor used for products or policies, so foundations like ownership, standards, metadata, change control, and feedback are treated as optional.

Without this discipline, knowledge stays unstable, hurting consistency, trust, and automation readiness. The fix is to apply product-grade governance to knowledge—clear ownership, scheduled reviews tied to releases, metadata standards, and auditable change control—so content is reliable for human decisions and safe automation.

GENAI AS A FORCE MULTIPLIER

GenAI becomes a true force multiplier only when it is anchored in governed knowledge. Governed knowledge is the body of service content—policies, procedures, FAQs, playbooks, and decision trees—managed with clear ownership, standards, version control, and audit trails.

On its own, GenAI can be inconsistent. Grounded in approved content, it improves both speed and quality. In practice, GenAI pulls context from CRM history and entitlements, detects intent and sentiment, drafts in-brand responses, and converts guidance into action—such as pre-filled forms, routed tasks, or scheduled follow-ups. Common high-value applications include:

- ✓ **Sentiment and intent detection:** Identify frustration or confusion, adjust tone, and escalate earlier when needed.
- ✓ **Next best step:** Suggest actions that reflect current status, policy limits, and entitlements.
- ✓ **Agent assist:** Retrieve the right article, generate a first draft, and display sources for quick confirmation.
- ✓ **Self-service:** Resolve routine queries with guardrails and hand off smoothly when confidence is low.
- ✓ **Decision transparency:** Explain the basis of outcomes to build trust rather than issuing blunt denials.

Safeguards keep results reliable. Retrieval-augmented generation ties answers to approved content. Privacy filters remove sensitive data before generation. Confidence thresholds trigger human review when uncertainty is high. Evaluation harnesses test AI outputs for accuracy, consistency, and policy adherence, and ties those results to KPIs like first-contact resolution, handle time, and CSAT—making changes to prompts, retrieval, and models auditable.

Measured results from our delivery programs include:

- ✓ Email Assist: ~23–25% monthly reduction in person-hours.
- ✓ Chat Assist: ~20–22% monthly reduction in person-hours.
- ✓ Conversational knowledge interface: ~30–32% reduction in search and drafting time.

"When outdated or contradictory knowledge feeds bots or GenAI models, errors scale up quickly. The core challenge is scale and fragmentation. Many organizations do not manage knowledge with the same rigor used for products or policies, so foundations like ownership, standards, metadata, change control, and feedback are treated as optional."

- ✓ AI-driven voice auditing: 100% call coverage at less than one cent per call.
- ✓ Product and sentiment analytics: ~15–17% reduction in manual analysis workload.

Each outcome connects to a clear mechanism—faster retrieval, better first drafts, full-volume quality assurance, or targeted analytics—so teams see how improvements are produced and can repeat them.

TRUSTED KNOWLEDGE + GENAI IN ACTION

When curated, trusted knowledge meets GenAI, customers experience faster, clearer service, and leaders gain traceability. Curated articles, policies, and decision paths form a single source of truth. GenAI uses that foundation to interpret intent, retrieve the right guidance, draft in-brand responses, and trigger the next best action. Agents spend less time searching and more time applying judgment.

Below are anonymized Straive client case studies, labeled by industry and context.

EdTech (North America).

Seasonal enrollment spikes created variability in quality and handle time. Centralizing the knowledgebase and standardizing playbooks across phone, chat, email, and web forms—then layering GenAI for retrieval and first-draft assistance—enabled the operation to handle ~138,000 calls, ~70,000 cases, and ~8,000 chats per year, maintain quality above 97% and CSAT at 4.5/5.0, improve phone AHT by 10–15% year over year, and achieve 98%+ accuracy in digital licensing and order entry. Governance kept the truth current; GenAI made that truth fast to apply.

Global Publishing (Multi-division).

Fragmented practices across K-12, higher education, research, consumer, and workforce-skills lines produced inconsistent experiences. With a unified CRM, 24×7 coverage, and GenAI-assisted drafting grounded in governed knowledge, the operation supported ~67,000 order entries and ~360,000 service interactions annually, transitioned ~150 roles offshore (about 60% savings), and reduced infrastructure costs by ~25%. Standardized content kept regional and product nuances accurate and visible across channels.

Professional Services (IT Helpdesk).

A multi-channel, fully managed desk adopted standardized SOPs and a governed knowledgebase, with GenAI assisting triage and drafting. The client realized a ~35% reduction in support costs, faster time to resolution, and redeployment of senior engineers to modernization initiatives. A unified 24×7 model serving 4,000+ users resolved most issues within 24 hours, supported by a self-help repository and full-coverage call auditing that strengthened coaching without expanding the quality team.

K-12 Learning (GCC-led service delivery).

Inside a Straive Global Capability Center (GCC), a Customer Service Delivery Pillar combined CX operations with IT services and engineering-adjacent work. The center processed 550,000–600,000 transactions per year and pursued 40–50% efficiency gains over three years by standardizing knowledge, embedding guardrails, and instrumenting quality audits. The GCC "factory" model—taxonomy and tagging, evaluation and release management, and prompt/retrieval stewardship—prevented drift as the program scaled across geographies and channels.

Across these examples, the pattern is consistent: curated, trusted data plus GenAI leads to optimized CX. Customers get faster resolution and clearer explanations grounded in policy; agents focus on higher-value work; and leaders gain a durable operating model. GCCs make the approach scalable by consolidating domain expertise, AI engineering, and knowledge operations so models are trained and evaluated on context-rich, approved data, and improvements persist as programs grow.

ROADMAP FOR ORGANIZATIONS

CX is, at its core, about people. AI-powered assistance eliminates repetitive, low-value work so frontline teams can concentrate on resolving complex issues. The result is empowered agents, an elevated role for the service function, and measurably better customer outcomes.

A successful path to GenAI-enabled service follows a disciplined sequence: build the foundations, validate impact in a controlled environment, and then scale with strong governance. The roadmap below is designed to help teams move quickly and confidently through each stage.

1. Create one source of truth.

Inventory repositories, merge or retire duplicates, and standardize content for omnichannel reuse. Add metadata (region, product, entitlement, policy version, effective dates). Make ownership and timestamps visible so freshness is clear. Shared, governed knowledge improves accuracy and ensures updates propagate everywhere.

"When curated, trusted knowledge meets GenAI, customers experience faster, clearer service, and leaders gain traceability. Curated articles, policies, and decision paths form a single source of truth. GenAI uses that foundation to interpret intent, retrieve the right guidance, draft in-brand responses, and trigger the next best action."

2. Define quality standards.

Set clear editorial rules for clarity, tone, and structure. Use templates for common intents and decision trees for complex paths. Schedule SME reviews aligned to release cycles; maintain audit trails; set measurable targets for accuracy and coverage.

3. Pilot high-volume journeys.

Choose a policy-defined flow (e.g., refunds, order status, password resets). Ground GenAI responses in approved content, apply privacy filters and confidence thresholds, and escalate uncertain cases to humans. Measure FCR, handle time, re-contacts, and effort; use findings to remove contradictions and close content gaps.

4. Expand to customer-facing channels.

Promote validated intents to self-service with the same guardrails. Low-confidence answers should hand off smoothly with full context. Track deflection and re-contact to confirm quality improves rather than simply shifting volume.

5. Build for scale.

Stand up governance (GCC or central team) to steward taxonomy, tagging, evaluation harnesses, prompt/retrieval patterns, and change logs. Pilot inexpensively; scale only approaches that meet accuracy and compliance thresholds. Balance global standards with local policy nuance.

6. Link metrics to content.

Tie CSAT, handle time, deflection, accuracy, and compliance to the specific articles and workflows used. Track accuracy at article and intent levels. Publish refresh cadences aligned to product and policy updates so improvements are continuous and auditable.

7. Prepare people and processes.

Train agents and SMEs to interpret AI suggestions, override when needed, and submit corrections that become permanent improvements. Recognize contributions and communicate clearly about changes and safeguards.

8. Plan the economics.

Model savings and reinvestment upfront (channel shift, deflection, handle-time gains, accuracy). Align GCC location and skill strategy to workload. Publish transparent benefit tracking for finance and operations.

Early success (90–180 days).

Agents and early self-service run on governed knowledge; agent-assist operates reliably in one or two journeys; standards and guardrails are active; accuracy is tracked; and a prioritized backlog drives ongoing improvements—creating a self-reinforcing cycle: govern, pilot, expand, refine.

LOOKING AHEAD

Customer service is moving to a paired model of GenAI plus agentic AI. GenAI interprets requests and drafts responses; an agentic layer executes policy-bound steps—checking entitlements, updating orders, scheduling callbacks—and closes the loop. This approach works only when actions are grounded in governed, auditable knowledge.

Orchestration will tighten. Teams will maintain a tool catalog (APIs, forms, workflows) and enforce rules through a policy engine. Operations will become event-driven: signals such as failed logins or delivery delays will trigger proactive actions. An evaluation harness will score accuracy, policy adherence, latency, and customer effort, and feed improvements into knowledge coverage, prompts, and processes.

Safety and scale will advance together. Programs will adopt data minimization, decision traceability, and tiered human-in-the-loop for sensitive steps. GCCs will centralize taxonomy, AI engineering, and governance—upholding global standards while accommodating local nuance. As predictive intent models improve, service will shift from reactive fixes to preventive care, amplifying empathy with automation so customers get the right answer—and action—the first time.

To translate insight into impact, organizations should first assess the quality of their knowledgebase and their readiness for GenAI. Straive's experts can benchmark current practices, identify gaps, and design a transformation roadmap.

To get started, connect with our experts to explore how governed knowledge combined with GenAI can elevate the customer experience.

About Straive

Straive is a leading data analytics and AI operationalization company that integrates cutting-edge solutions into core workflows to enhance efficiency, elevate user experience, and drive revenue. Serving industries such as Pharma & Life Sciences, Logistics, Publishing, EdTech, BFSI, and more, Straive supports clients in over 30 countries. With a globally distributed talent pool across the Philippines, India, the U.S., Nicaragua, Vietnam, the U.K., and its headquarters in Singapore, Straive delivers transformative, scalable outcomes worldwide.

Website: http://www.straive.com

Contact: contact@straive.com