



EXECUTIVE POINT OF VIEW

The Utility Customer Experience Playbook for Agentic AI Success

Ground truth, guardrails, and the economics of selective automation in M2C

EXECUTIVE ABSTRACT

Utility companies face intense pressure to digitize, yet shifting customer interactions entirely to automated bots poses immense strategic risk. True compassion cannot simply be programmed. It must be grounded in actionable operational data. The same is true of governance: an autonomous agent operating on live customer accounts must be auditable, compliant, and defensible, and none of that comes from the model. It comes from what the utility knows about its own interactions.

This paper argues that granular interaction insight is the essential foundation for autonomous agentic AI in meter-to-cash, serving as the source of both empathy and guardrails. It sets out how utilities build that foundation, where automation is safe and where it must never go, and what early trials show about the economics of getting that calibration right.



1. The Trap of Automation Without Context

The Meter-to-Cash (M2C) lifecycle is complex, highly regulated, and deeply critical to daily life. Its failures, such as a \$600 disputed bill, a winter freeze outage, or a disconnection notice arriving in the coldest week of the year, are not routine service events. These are high-stress moments in which a customer’s trust in their utility is either confirmed or lost.

Rushing to virtual assistance in that environment creates generic, rule-based systems that deflect calls without contextual or demographic awareness. For customers less comfortable with text-based bots, or for anyone in genuine crisis, being met with “press 1 for bill pay” is not merely unhelpful. It is a breach of trust. The wariness this creates is not anecdotal: in a 2024 Gartner survey of more than 5,700 customers, 64% said they would rather a company not use AI in customer service at all, and the leading concern was losing the ability to reach a human being.

The difficulty is that the interactions requiring human judgment do not announce themselves. They arrive through the same channels, in the same words, and look identical to a routing engine.

Straightforward to Automate		Looks Routine, Needs a Human	
	Meter read verification	“Why is my bill so high?” three weeks after an AMI rollout	
	Move-in / move-out scheduling	A payment-plan question that is an early hardship signal	
	Balance and due-date inquiry	An outage call from a premise on life-sustaining equipment	
	Field appointment confirmation	A move-out request on an account already under disconnection notice	

2. Why CX Insights Have to Come Before Autonomous AI

Most utility executives are aware of Agentic AI, and many are already exploring it today. The promise of intelligent agents taking full ownership of complex tasks, such as resolving a billing dispute end to end, rescheduling a field visit, or arranging a payment plan, autonomously with minimal human involvement, is huge, and it is no longer theoretical. Gartner projects that by 2029, agentic AI will autonomously resolve 80% of common customer-service issues without human intervention, driving a 30% reduction in operational costs — a scale of change few utility organizations have fully prepared for.

But jumping straight to autonomous agents without first mastering operational customer data is like putting an F1 car on autopilot before mapping the track. The car is not the risk. The unmapped track is. And in a regulated utility, that track includes commission scrutiny, disconnection moratoria, medical baseline registries, and tariff structures that shift with the season.

Utility interactions are not simple. Advanced AI models analyze them at scale to break them down into deep, granular sub-categories, for example:

Surface Category:

Billing Issue

Sub-Category:

High Bill Variance Post-AMI Rollout

Root Cause:

Estimated reading error during summer peak tariff transition

The distinction matters enormously. An agent operating at the surface category closes a billing complaint. An agent operating at the root cause level surfaces a meter data quality issue affecting an entire rollout cohort: a different action, a different remedy, and a very different regulatory exposure when the same decision is repeated ten thousand times autonomously.

You cannot govern what you have not categorized.

3. CX Insights:

The Compass for Compassionate AI

Virtual assistance is an operational and behavioral challenge, not an engineering problem. The utilities that succeed with it will not be those that select the best model. They will be those that understand their own interactions well enough to know what a model should be permitted to do. That understanding comes from a Ground Truth layer: millions of categorized interactions, classified through domain expertise in M2C business processes along two axes at once, namely interaction intent and customer emotional state. Both are required. Intent alone will route a distressed customer efficiently into precisely the wrong experience.

A mature Ground Truth layer answers two vital questions:

- **Which processes can be perfectly automated for efficiency?**
(Standard billing inquiry, meter read verification.)
- **Which interactions require a human-in-the-loop (HIL)?**
(Hardship support, emergency services.)

Insights also reveal that empathy requirements are highly demographic-specific. While a millennial may prefer a bot for simple billing, a senior customer may find that same bot confusing and distressing.

The same interaction type, handled identically, produces opposite outcomes.

Insights dictate orchestration.



4. The Architecture of Empathy

What does an empathetic, compassionate AI built on CX insights actually look like in operation?


It looks like an intelligent, technology-agnostic orchestration layer that works alongside the systems a utility already runs, including Oracle CC&B and Customer Cloud Service, SAP, and existing CRM and field service platforms, without vendor lock-in. Rather than replacing the system of record, it governs how interactions are routed across it.

Using CX insights to calibrate, that layer deploys specialized micro-agents against defined M2C functions. Implementation of an empathetic **‘Utility Buddy’** showcases this agentic model:



Payments Agent

Frictionless, secure, insight-optimized payment journeys.



Billing Concierge Agent

Complex M2C billing questions and deferred payment structures.



Outage Communications Agent

Automated outage updates that keep grids and customers connected.



Field Service Agent

Rescheduling field appointments with speed and precision.



Service Transfer Agent

Automated end-to-end move-in and move-out workflows.

Each agent operates strictly within the boundaries the Ground Truth layer defines and, critically, hands off the moment an interaction crosses one of them.

5. From Insights to Guardrails: Governance, Trust, and Liability

Categorized CX data does something beyond routing. It shows AI agents the exact decision trees and compliance rules used by successful human agents, which is what converts a customer experience initiative into a governable enterprise system.

These are the questions the CIO, CFO, or General Counsel must ask before signing off on autonomy. Each is answered by data, not by model selection.

Concern	What Interaction Insight Provides
Auditability	Every autonomous action is traceable to a category, a rule, and documented human precedent — answerable to a commission complaint or a discovery request.
Regulatory conformance	State-specific disconnection procedure, moratorium windows, medical baseline and life-support registries, and LIHEAP handling encoded as agent guardrails rather than inferred by a model.
Delegated authority	Where agents adjust bill credits, initiate service holds, or update preferences: defined action limits, value thresholds, and mandatory escalation points.
Production readiness	Taxonomy coverage of live interaction volume before launch; containment, escalation accuracy, false-automation rate, and drift monitoring after it.


Insight is not the training data for autonomy. It is the control layer.

6. The Critical Intersections: The Human-First Safety Net

The core principle of any credible orchestration model is knowing when NOT to deploy a micro-agent. Certain trigger scenarios must bypass automation entirely for immediate human interaction:

- **Emergency / Safety.** Reports of gas leaks, downed wires, or life-sustaining medical equipment failures require instant human validation and empathetic assurance.
- **Hardship / Vulnerable Customer Support.** Disconnection notices, LIHEAP requests, or distressed payment arrangements flag immediately to empathetic CSRs trained in hardship procedures.
- **Demographic & Age-Group Preference.** Data models map successful interactions by age, enabling intelligent prediction of when a customer will find a bot frustrating.

One caveat deserves stating plainly: an immediate-bypass path presumes 24×7×365 human availability. Utilities should treat that as a staffing commitment rather than an assumption. A bypass route that lands in an after-hours queue is worse than no bypass at all, because it manufactures a documented failure at exactly the moment that matters most.



7. What the Trials Show:

The Economics of Selective Automation

Calibrating automation levels based on call type and customer sentiment is gaining traction in early implementations of agentic AI in utilities, and is emerging as the standard for achieving reliable outcomes. These early deployments demonstrate substantial cost savings over legacy models, reducing operational risk while maintaining top-tier service.

 Call Type / Scenario	 Cost Savings	 Business Benefits
Standard bill pay and meter verification	Up to 85% reduction	Frictionless self-service and improved satisfaction
Simple move-in / move-out	Up to 75% reduction	Shifts volume off-peak, 24x7
High-bill variance & tariff explanation	Up to 50% reduction	Cuts handle time, retains judgment
Payment arrangements & hardship (LIHEAP)	Up to 40% reduction	Preserves CSR capacity for resolution
Emergency gas leaks / downed wires	Risk avoidance	Protects highest-stakes response

The argument these numbers carry extends far beyond cost control or operational efficiency. Calibration is what protects customer trust and elevates the overall service experience. Automating routine workflows frees skilled human capacity for complex, high-stress interactions that demand genuine care, simultaneously lowering cost to serve and response times. Efficiency and empathy are not a trade-off here. They are the same decision.

8. From Passive Q&A to Proactive Strategy: The Utility-Specific LLM

Generic LLMs fall flat on complex M2C tariffs and regulatory compliance. By applying disciplined prompt engineering to domain-trained models, the technology shifts from reactive handling into a proactive intelligence engine, surfacing strategic insight directly to CXOs:

Proactive Insight: “We detected a 34% spike in outage-related call escalations in Zone B. Primary driver: conflicting restoration times on web portal vs. IVR.”

Recommended Action: “Sync portal updates with field technician dispatches and trigger proactive SMS alerts to affected accounts to suppress call volume.”

Interaction data then feeds back into the model, sharpening performance and uncovering the next efficiency gap. The foundation does not simply enable autonomy once. It compounds.



Conclusion: CX Insights as the True North of Virtual Assistance

The next generation of Utility CX will not be defined by chatbot speed, but by the empathy of its context.

That context is built, not bought. It comes from categorizing interactions to root cause, classifying emotional state alongside intent, encoding compliance rules from documented human practice, and deciding deliberately where delegated authority ends. It is unglamorous work, and it is what determines whether an autonomous agent can be audited, defended before a commission, and trusted on a live account at two in the morning during a storm.

Utilities that establish this foundation will automate more confidently than their peers, not less. They will know precisely which interaction requires rapid AI action and which demands the irreplaceable power of a human voice.

Sutherland is an AI-driven business transformation partner to enterprises, including several Fortune 500 clients. We design, run, and automate business operations at scale, helping clients improve efficiency, strengthen customer experiences, and deliver measurable outcomes.

Most companies either build AI-led capabilities or run operations. Sutherland does both. We run complex, mission-critical processes while engineering the AI-led solutions that sit within them. This closed loop helps clients move beyond AI pilots to reliable, governed outcomes that improve and scale over time.

With deep operational expertise, advanced products, platforms, and engineering, we help clients modernize operations and build future-ready businesses.

