

# Lean → LLM Readiness Checklist (Contact Center Edition)

Version: 1.0 | Date: 2026-03-02 | Use case: Evaluation checklist for CX Ops leaders considering agent assist / GenAI in a contact center.

This checklist helps you validate whether your service workflows are stable enough for LLM-based assistance to scale value without amplifying variation. It focuses on operating-model control points, governance loops, and process-stability signals (not tool features).

## A. Process stability & decision boundaries

- Top contact drivers are defined and stable (taxonomy is consistent across queues, teams, and shifts).
- A clear 'definition of done' exists for each top driver (what a correct resolution looks like).
- Decision boundaries are documented: what can be resolved in-line vs. what must escalate.
- Exception categories are defined and owned (who reviews, updates, and approves changes).
- Escalation triggers are explicit and testable (not dependent on individual agent judgment).

## B. Knowledge & input quality

- A clear source-of-truth hierarchy exists for knowledge (which docs override which).
- Knowledge articles and scripts have owners, review cadence, and version control.
- Required intake fields exist to reduce ambiguity (e.g., state/context fields for top drivers).
- Duplicate or conflicting macros/scripts are reduced to 'golden paths' for top scenarios.

## C. QA & governance loops

- QA scorecards map to measurable behaviors (not subjective interpretation).
- QA calibration cadence is defined (auditors align routinely; disputes have a resolution path).
- Feedback loops exist to update taxonomy, knowledge, and exception rules based on real cases.

## D. Measurement signals (prove stability, not just adoption)

- You track reopen / repeat contact trends by top driver (not just overall).
- You monitor handle time variance (not only average) across cohorts (new vs. tenured, teams, shifts).
- You monitor escalation accuracy (right issues routed to the right tier early).
- You monitor 'reason changed after first tag' events as a taxonomy health signal.
- Knowledge freshness cadence is visible (what was updated, when, and why).

How to use this: Start with your top 10–20 contact reasons. Baseline stability signals (variance, reopens, tag changes). Stabilize the workflow and governance loops first, then expand AI scope only when outcomes become repeatable.

Note: This checklist supports evaluation and readiness validation. It does not imply guaranteed outcomes and does not prescribe vendor-specific implementation steps.