

Seven questions to ask before you believe an AI resolution rate.

For support, CX and finance leaders reading a vendor case study, a business case or their own dashboard. Every headline rate leaves some of these blank. Fill them in before you set a target or a headcount.

01	STARTING POINT	How much did customers already solve on their own before AI? A company with little self-service sends easy questions to support, so its AI posts a big number. Strong self-service leaves only the hard leftovers.
02	MIX	What share of contacts are simple, and what share need judgment? B2B, complex products and regulated work carry more judgment. A rate earned on one mix doesn’t transfer to another.
03	DEFINITION	What counts as “resolved”? “Ended without a hand-off” is not “the customer confirms it’s solved.” Ask whether abandonment counts as unresolved, and what repeat-contact window applies.
04	DENOMINATOR	Is it a share of tickets, or a share of all customer demand? Restated against total demand, a lower headline can be the better result. See the example below.
05	WORK	How many hours did it remove, net of the new work it created? Easy tickets are short, so a share of tickets overstates a share of hours. Knowledge, quality review and escalation design are new work.
06	COST	What does each resolved issue cost, all in? Per-resolution, per-conversation and per-seat pricing reward different numbers. Divide the full cost by issues actually resolved.
07	QUALITY	How do AI-resolved issues compare with human-resolved ones? Repeat contact, satisfaction and errors, side by side. A rate without quality is half a result.

Why the denominator matters	GROUND FLOOR	MATURE TEAM
Two companies, the same AI. The one that built self-service years ago posts the lower headline and solves more for its customers. An illustration, not a benchmark.	Customer questions a month	100
	Solved by self-service first	0
	Reach support as tickets	100
	Resolved by AI	40 (harder)
	Vendor headline	65
	Questions solved without a person	15
	65%	38%
	65%	75%