

The CFO's Guide to Contact Center ROI

How to quantify a return you can defend to the board, and stop funding the contact center on faith.

THE 40-SECOND VERSION:

You can measure what your contact center costs to the dollar. You probably can't measure what it returns to within a million. That asymmetry is why it keeps getting cut. Modern cloud contact centers close it, every second of handle time, every avoided repeat contact, every point of service level maps to a number in cost, risk, or revenue. This guide shows how the math is built, what the benchmarks look like, and how to walk into a board meeting with a defensible ROI band instead of a feature list.

The Numbers, Up Front

Financial outcome	Reported benchmark
Payback period	~23 months (G2-reported)
Operational cost reduction	30–50% lower operating costs
Cost per contact	\$2–\$8 decrease per contact
Total cost of ownership	40–60% lower TCO over three years
Time to value	~2 months average implementation (G2)
Agent turnover	Up to 50% reduction

Where these come from: payback (~23 months) and implementation time (~2 months) are from UJET's third-party G2 profile. The rest, operating-cost reduction, cost-per-contact savings, TCO reduction, and agent-turnover reduction, are UJET business-case benchmarks: modeled reference ranges from UJET's own analysis, backed by the customer results in this guide (a major US bank's 10% cost reduction, a streaming provider's 23% AHT cut, Capital on Tap's 92% SLA, Zettle's halved training). Reference ranges, not guarantees; UJET-authored analysis, not independent measurement.

Your own defensible number is built in Chapter 4.

Contents

(01)	The Asymmetry That Keeps the Contact Center Underfunded	→
(02)	The Three Things You Already Fund	→
(03)	What the Numbers Look Like	→
(04)	How the Math Is Built, the 100-Agent Model	→
(05)	Proof: The Levers Working in Real Operations	→
(06)	Finance's Hardest Questions, Answered	→
(07)	Build Your Number with Us	→

(01) The Asymmetry That Keeps the Contact Center Underfunded

Every CFO funds what can be measured, and the contact center has always been half-measurable. The cost side is precise to the penny, headcount, licenses, telecom, seats. The return side is written in a different language entirely: average handle time, first-contact resolution, service level, CSAT. The people who run the function speak one; the people who fund it speak the other. The two rarely meet in the same sentence, so a function that touches nearly every customer gets managed as overhead to be trimmed.

That was an accounting choice, and it has quietly become the wrong one. Modern cloud contact center platforms made the link between operational performance and financial outcome direct and traceable. A second saved on handle time is a quantity of agent hours you can hand back to the business. A point of resolution quality is fewer repeat contacts, and fewer repeat contacts is cash the company keeps. The category has a name now, the **Experience Center**: a contact center run as a revenue engine rather than a cost line. This guide is about making that return explicit enough to approve, or to pressure-test, on financial logic.

Finance's caution was never about undervaluing customers. It was about three kinds of uncertainty, and uncertainty is the one thing a budget owner can't easily sign. The first is **legacy technology debt**, upgrading on-premise systems is expensive and unpredictable, and the bill rarely stops at the license. The second is **hidden, unpredictable cost**, pricing loaded with usage-based charges and

per-API fees that make next year's number impossible to forecast. The third is **vendor lock-in**, commitments that can't flex with the business or ride the cloud spend you've already agreed to.

Those three are not bad luck. They are symptoms of an architecture problem: a CCaaS stack stitched together from add-ons, each with its own meter. The fix is architectural. A platform built natively on Google Cloud, purchased through the Google Cloud Marketplace against an existing cloud commitment, with a CRM-first design and no customer PII stored on the vendor's servers, removes the usage surprises, the lock-in, and a category of data risk at the same time. Hold that thought, it reappears in Chapter 6, where finance's hardest questions get answered.

(02) The Three Things You Already Fund

A contact center investment gets approved when it moves at least one of three things you already track. Almost every credible business case is some combination of them.

Cost to serve, the operating expense behind every interaction, the staffing required to handle volume, and the total cost of ownership across a sprawl of platforms and add-ons. **Risk and resilience**, the exposure from downtime, compliance gaps, and preventable churn; an outage means lost revenue and reputational damage on top of the operational hit. **Revenue impact**, retention, conversion, upsell, and the churn that leaks out when service slips on high-value accounts.

Value arrives through six specific levers, and each maps cleanly back to one of those three:



Handle time falls as after-call work is automated away → cost.



Repeat contacts fall as resolution quality improves → cost.



Service levels rise, better SLA attainment, lower abandonment, faster answer → risk and revenue.



Volume shifts to lower-cost digital and self-service channels → cost.



Agent ramp and turnover drop as training time shrinks → cost.



Total cost of ownership falls as tools consolidate onto one platform → cost.

Each lever is measurable. That's the whole game: measurable levers, mapped to the three dimensions finance already reports.

(03) What the Numbers Look Like

Before modeling your own environment, it helps to know the outer envelope of what modern CCaaS platforms have returned. The table at the front of this guide holds the reference ranges, a G2-reported payback of roughly 23 months and ~2-month implementations, plus 30–50% lower operating costs, \$2–\$8 off cost per contact, 40–60% lower TCO over three years, and up to 50% lower agent turnover.

No CFO signs off on a range. The point of the benchmarks is narrower and more useful than that: they establish that this investment operates at a scale worth modeling precisely. A 40–60% TCO swing on a multi-million-dollar cost base is not a rounding error. It is a number worth an afternoon of finance's time, which is exactly what the next chapter is for.

(04) How the Math Is Built, the 100-Agent Model

The model has two halves. On one side, the inputs you already own: monthly contact volume across calls, chats, and messages; baseline handle time and after-call work; fully-loaded hourly cost or cost per FTE; staffing and occupancy assumptions; current platform and add-on costs; and your target service and CSAT goals. On the other, the outputs finance needs to report: annual OpEx savings, payback period, a 12–36 month ROI range across best, base, and worst cases, and a sensitivity view of which assumptions move the result most.

To see inputs become outputs, take a representative 100-agent contact center: roughly 475,000 contacts a month, 300-second average handle time, a \$16 hourly agent rate, 85% occupancy, 60% first-contact resolution. These are illustrative figures chosen to show the mechanics, not a specific customer, and not a promise.

Working the levers, handle time falls by 21 seconds per contact and first-contact resolution improves. Together those avoid about 4,752 repeat contacts a month and shift more than 10,000 interactions into self-service. The staffing math follows: the same volume is served by roughly 66 agents instead of 100, 34 fewer FTEs, and total monthly cost drops by about \$181,000.

What it adds up to. In this illustrative model, the changes produce roughly \$2.2M in annual savings, about 38% of total contact center cost. The cleanest single calculation sits inside it: 21 seconds saved across ~138,000 agent-handled contacts a month is over 800 agent-hours returned to the business every month.

The number that matters isn't this one. It's yours. And you don't have to wait for a sales call to see a first version of it.

(05) Proof: The Levers Working in Real Operations

Benchmarks establish the range; customer outcomes show the levers working. Each result below pairs a measured outcome with the way finance would translate it into dollars.



A major US regional bank, cost and digital shift.

Modernized its contact center, moved routine service into digital channels, and landed a 15% reduction in call volume, a 10% reduction in contact center costs, and a 50% increase in digital chat. *For finance:* the call-volume reduction applies directly against fully-loaded cost per call; the channel shift lowers cost per resolution and eases staffing pressure.



Capital on Tap, service level and CX. This fast-growing fintech consolidated onto a modern platform with intelligent routing, lifted SLA attainment to 92%, cut average hold time by 12%, raised CSAT from 4.4 to 4.6, and now answers 90% of calls in under 20 seconds. *For finance:* SLA and hold-time gains reduce churn risk on at-risk accounts and make operations more predictable.



EverWash

EverWash, speed and efficiency. This subscription car-wash network cut average response time by 70%, from more than 60 seconds to 18. *For finance:* faster response reduces abandonment and raises capacity per agent hour. Model the freed capacity as deflected headcount or as volume growth absorbed without hiring.



Zettle by PayPal, onboarding cost. This global payments provider adopted AI-guided onboarding, raised CSAT 10%, and halved agent training time, at a lower cost than its previous solution. *For finance:* cutting ramp time in half reduces onboarding cost with the very next agent cohort.



A leading streaming service, productivity. Rolled out a unified agent desktop and AI assist across high daily volume and cut average handle time by 23%. *For finance:* AHT reduction times volume times fully-loaded cost is the cleanest calculation there is. At scale, a 23% cut converts straight into recovered agent hours.

The same levers hold across service patterns. Turo improved SLA by 91%, cut AHT by 28 seconds, and lifted CSAT 5%. Wag! improved SLA by 17%, halved average wait time, and dropped call abandonment by 7%. A blinded grocery-delivery company resolved contacts 30% faster, cut repeat contact by 5.6%, and improved SLA by 18% on voice and 23% on chat. A blinded mobility platform improved SLA by 61% over three years while scaling from 40 to 242 agents and routing 65% fewer calls to live agents.

(06) Finance's Hardest Questions, Answered

A business case survives its own review. Three objections come up in nearly every finance conversation. Here are the answers.

"I don't trust operational metrics."

Convert them. Every operational number ties to a financial outcome: handle time becomes recovered agent hours, service level becomes retained revenue on at-risk accounts, resolution quality becomes avoided repeat-contact cost. Pair that with fixed before-and-after measurement windows and a set reporting cadence so the baseline can't quietly drift, and the metrics become financial evidence. One caution on the baseline itself: a QA sample of 2% of conversations is a guess dressed as a number. Conversational analytics that read 100% of interactions, the job **Spiral** does, give you a baseline you can defend.

"What about hidden costs?"

Put them all on the table before comparing. A complete TCO view includes implementation and professional services, telecom and carrier charges, analytics, QA and workforce-management add-ons, integration maintenance, and the admin and IT overhead the platform demands. This is where architecture pays the CFO directly: transparent, marketplace-based pricing, purchasing through the Google Cloud Marketplace against an existing cloud commitment, makes the total predictable and removes the usage-based surprises. A CRM-first platform with no customer PII stored on the vendor's servers also takes a class of compliance and data-breach exposure off your risk register.

"What if adoption is low?"

Adoption is a governance question, so govern it. Define the KPIs up front, the share of interactions using AI assist, the share of contacts shifted to self-service, the share of agents consistently using integrated workflows, and attach a 30/60/90 adoption plan with executive checkpoints. Value realization becomes something tracked on a schedule rather than assumed at signing.

(07) Build Your Number with Us

This guide gives you the framing, the benchmarks, the mechanics, and the proof. The one thing it deliberately doesn't hand you is a false-precision spreadsheet, because the number that earns a board sign-off is built from your real operation.

Two ways to get there:

- 1 Run the interactive ROI estimate.** Drop in your volume, handle time, agent cost, and current tool spend, and see a first-pass savings and payback figure in a few minutes. It's directional, self-serve, and enough to decide whether this is worth a deeper look. Start at ujet.cx/roi
- 2 Book an ROI review.** A UJET sales engineer builds the full model with you on your real baseline, annual OpEx savings, payback period, and an ROI band across best, base, and worst cases you can defend to the board line by line.

Bring the metrics you already track, monthly volume, baseline handle time, fully-loaded agent cost, current platform and tool spend, and your target service levels. That's enough to build a first-pass case in a single session.

NEXT STEP

The next step. Start with the calculator for a number in minutes. Book the ROI review when you want one you can take to the board. Either way, you leave with a figure you can defend.

Start here: ujet.cx/resources/tools/roi