Revenue-Growth.AI

CLIENT CASE STUDY

Confidential — For Executive Review

AskNicely × Revenue-Growth.ai

AI-Native Deal Intelligence — Deployed In-Tenant, Live in 14 Days

CLIENT	SEGMENT	ARCHITECTURE	INVESTMENT	DEPLOYED
AskNicely	Mid-Market B2B SaaS	100% In-Tenant	$15K–$20K One-Time	2024

TL;DR Executive Summary

Three numbers. That is the entire investment thesis.

+92%
WIN RATE INCREASE
From baseline — not incremental lift

−50%
SALES CYCLE REDUCTION
Deals closed in half the prior time

3mo→1mo
AE RAMP TIME
−67% time to first close

Day 15

Payback period on the full CapEx investment. A single month of the fractional RevOps retainer AskNicely was evaluating costs more than the entire Revenue-Growth.ai deployment. The infrastructure paid for itself before the calendar turned.

01 The Financial Pain — The "Before" State

AskNicely's revenue team was not failing. They were bleeding — slowly, predictably, and invisibly.

Their Franken-Stack was operating exactly as designed: Salesforce as a passive archive, reps manually translating call notes into CRM fields, leadership pulling fragmented pipeline reports and reconciling them by hand before every forecast call. The data existed. It just lived in the wrong places, arrived too late, and required human interpretation at every step.

THE QUANTIFIABLE BURN

40% CAPACITY	**RevOps bandwidth consumed by manual CRM entry**, pipeline interrogation, and deal review prep — not revenue-generating activity.
$450–$600	**Per user per month in tool sprawl** across Salesforce, Gong, Outreach, and ZoomInfo. Each tool generated more data. None of it talked to the others without manual intervention.
3-MO RAMP	**Every new AE started from zero**, relying on tribal knowledge passed through Slack threads and informal mentorship — not institutionalized deal intelligence.
NO SIGNAL	**Inconsistent win rates driven by rep-to-rep variance.** The company had a defined value framework. Whether individual reps executed it on any given deal was entirely subjective.

The result was a forecasting function built on confidence intervals derived from gut feel. Leadership was making resource allocation decisions — headcount, pipeline coverage, board guidance — based on data that was stale the moment it was entered.

This is not an operations problem. It is a structural one. When reps are functioning as manual data routers between calls and CRM fields, they are not selling. They are performing data entry at an AE salary.

The alternative AskNicely was evaluating: A fractional RevOps retainer at $10,000–$20,000/month. Year 1 outlay: $145,000–$265,000. Expected time to value: 6–12 months — following a 4-week discovery audit before a single line of execution was delivered. The result would have been a set of process recommendations and a renewed dependency on external human consultants to implement them.

02 The Architecture & Deployment — The "How"

Revenue-Growth.ai deployed an AI-Native Deal Intelligence framework directly inside AskNicely's existing CRM environment. No new software. No new logins. No workflow changes for the sales team.

This distinction is load-bearing.

The standard objection from every CISO evaluating AI for revenue operations is: *"Our SOC 2 review will take six months — and they'll likely reject it."* That objection applies to every vendor that routes proprietary pipeline data through external servers to run LLM inference. It does not apply here.

> **SOC 2 bottleneck: bypassed architecturally, not contractually.** Because the Revenue-Growth.ai framework is constructed entirely within the client's existing CRM tenant, no data ever leaves the environment. The LLM processes deal data inside the same security boundary that Salesforce or HubSpot already operates within — a boundary that has already passed every compliance review the client has ever conducted. There are no NDA carve-outs, no data processing addendums, no third-party server agreements to review. The CISO's objection is resolved at the infrastructure layer before the procurement conversation begins.

THE 14-DAY DEPLOYMENT TIMELINE

DAYS 1–3

Architecture Mapping

CRM structure, sales playbooks, deal stages, buyer personas, and historical win/loss patterns mapped. The framework inherits the client's institutional context — not a generic template.

DAYS 4–10

In-Tenant Agent Construction

AI agents constructed natively inside the existing CRM using the client's own secure API endpoints. Every inference is hyper-contextual to AskNicely's specific buyer patterns and sales motion.

DAY 14

Autonomous Execution Live

Zero behavioral changes required. Agents begin analyzing calls, scoring deal health, mapping buyer signals, flagging pipeline risk, and generating AI-authored close plans — autonomously.

> **What "zero disruption" means in practice:** AskNicely's AEs never received a new login. There was no change management initiative. There was no training session. Reps sold. The infrastructure executed everything else.

---

03 The ROI & Business Impact — The "After" State

> **This was a $15,000–$20,000 one-time CapEx investment. The following outcomes were immediate.**

WIN RATE +92%

Objective, AI-driven deal analysis replaced subjective rep intuition as the primary input to deal strategy. The framework evaluated every active opportunity against AskNicely's historical win patterns in real time, surfacing precise next-step guidance calibrated to actual deal signals — not rep confidence.

The 92% win rate increase is not incremental lift over a prior AI baseline. It is movement from a rep-dependent, gut-feel process to a systematized intelligence layer that operates consistently across every deal in the pipeline.

SALES CYCLE LENGTH –50%

Pipeline bottlenecks — deals stalled in negotiation, stuck awaiting internal champion follow-through, or degrading due to missed signals — were identified and flagged before they became forecast casualties. The autonomous close plan system eliminated the lag between deal review meetings and next-step execution. Deals moved faster because the intelligence to move them existed at the moment of each CRM interaction, not three days later after a manager's pipeline review.

AE RAMP TIME 3 MONTHS → 1 MONTH

This is the metric that directly compounds. Every new AE who ramps in one month instead of three represents two additional months of quota-carrying capacity per hire. For a growing revenue team making multiple hires per year, this is a structural improvement to the unit economics of every future recruiting decision the company makes.

New AEs inherited AskNicely's institutional deal intelligence on Day 1. Historical win patterns, objection handling, buyer persona signals, and playbook adherence criteria were embedded into the CRM environment the rep was already operating in. There was no onboarding program to complete. The knowledge was in the system.

THE CAPEX MATH — FULL ALTERNATIVE COMPARISON

ALTERNATIVE	YEAR 1 COST	TIME TO VALUE	DATA SOVEREIGNTY
Fractional RevOps Retainer	$145,000–$265,000	6–12 months	Data leaves your environment
Full-Time RevOps Manager (fully loaded)	$120,000–$160,000	3–6 months to ramp	N/A
Revenue-Growth.ai	**$15,000–$20,000 (one-time)**	**14 days**	**100% in-tenant. You own it permanently.**

The CapEx investment did not create a dependency. AskNicely owns the framework permanently. There is no retainer to cancel, no platform subscription to renew, no external service that can be deprecated. The intelligence infrastructure operates inside their CRM as a permanent operational asset — with no ongoing cost structure attached to it.

04 The Verdict

> *"The impact on our win rates and time to close was immediate. After using Revenue-Growth.ai, I can't imagine helping lead a revenue team without it. It's not a tool — it's infrastructure."*
>
> **— Alex Burkholder, VP Revenue, AskNicely**

Alex Burkholder's characterization is precise, not rhetorical.

AskNicely was at the stage every mid-market B2B SaaS company eventually reaches: revenue motion too complex to run on rep intuition, but not large enough to absorb a six-figure annual RevOps consultancy spend. The Franken-Stack had been assembled. The data existed. The playbook existed. The intelligence layer to connect them did not.

Generic AI was not a viable option. Routing AskNicely's proprietary pipeline data to an external LLM would have triggered a compliance review that would have delayed any deployment by six months — and introduced an ongoing data governance liability with no clean resolution.

A fractional RevOps retainer would have produced a process map and a set of recommendations. It would not have produced autonomous deal intelligence operating inside the CRM at the deal level, on every active opportunity, in real time.

> The in-tenant architecture was the only solution that resolved all three constraints simultaneously: **speed of deployment, data sovereignty, and permanent operational ownership.** That is not a positioning claim. It is a structural reality.

A 92% win rate increase and a 50% reduction in sales cycle length were not the byproducts of better processes or harder-working reps. They were the direct output of replacing manual data routing with autonomous deal intelligence — and deploying that intelligence inside the environment the team was already operating in.

Day 15. The payback period was Day 15.
