

ROSHAN CHANDRASENA

ENTERPRISE AI COMMERCIALIZATION

I commercialize enterprise AI by turning customer evidence into better agents, better workflows, and repeatable revenue.

Canada-based | Canada/U.S. remote | Business travel | rosh@roshdc.com | roshdc.com

AI rarely fails commercially because the model is unimpressive. The harder problem is connecting buyer need, product behavior, workflow, ownership, economics, and adoption. My advantage is operating across that entire system: top-ranked commercial performance, live AI decision logic, enterprise pilots, product-GTM change, and independently built technical platforms. That range lets me diagnose the binding constraint and turn field evidence into a system the company can improve and scale.

Top-ranked during tenure across 40+ AEs | \$4M+ cash collected | 5 live AI applications | 1,000+ mature outbound attempts/day

EVIDENCE

I know where AI commercialization breaks because I have owned the customer outcome

At Planswell, I ranked at the top during my tenure across a 40+ AE organization, generated \$4M+ in cash collected, and produced roughly 1.5-3x the team-average cash collected per customer. That performance exposed where the wider system was leaving value behind. I helped redesign pricing, packages, custom solutions, qualification, handoffs, follow-up, support, and account-management requirements; as the changes spread, team-average customer value rose by roughly 40% and retention also improved.

I translate field evidence into product and operating decisions

Planswell developed five live conversational-AI applications spanning outbound prospecting, sales and nurture, adaptive seller training, self-service financial-plan guidance, and configurable advisor acquisition. I used calls, transcripts, conversation depth, meetings, handoff conversion, order size, and time to close to evaluate agent behavior, diagnose failure points, and guide changes to prompts, model configuration and weighting, decision trees, qualification, routing, live-transfer thresholds, scorecards, and release acceptance. Engineering owned production code and infrastructure; I helped lead customer-facing logic and commercial-readiness decisions. The operating loop was simple: measure behavior, change the system, and validate whether the commercial outcome improved.

I can carry the system from individual performance to enterprise adoption

That feedback loop also shaped enterprise execution. I secured and helped launch Planswell's first four enterprise engagements across financial services, coordinating executives, compliance, sales, marketing, operations, product, and implementation. The work connected a credible use case to workflow, adoption, governance, success measures, and pilot execution—so AI capability became an operating commitment rather than a demonstration.

INDEPENDENT PROOF

C2 Lattice makes agent coordination inspectable: a Python-standard-library control plane with 22 tools spanning task ownership, authority, dependencies, signed identity, file locks, recovery, and a live dashboard. The public source is supported by 465 named local tests across six suites, separate E2E coverage, 28/28 clean-checkout acceptance, and green hosted CI on Ubuntu/Python 3.12 and 3.13. Relay extends the same operating discipline into secure private AI workflows, grounded context, evaluations, model routing, and cost visibility; Area10 applies it to point-in-time economic intelligence.

WHAT THIS MEANS I help leadership teams turn promising AI into a commercial system: diagnose the constraint, win customer trust, evaluate how agents and workflows behave, translate evidence into product and operating changes, and build a repeatable path from deployment to measurable value.

PROOF [LinkedIn](#) | [GitHub](#) | [AI & ML ePortfolio](#)