

ROSHAN CHANDRASENA

ZERO-TO-ONE GTM & OPERATIONS

I turn early traction into a coordinated growth system that can learn, scale, and survive handoffs.

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

Growth usually breaks between functions, not inside one of them. Product, sales, delivery, customer value, and economics can each look reasonable while the system underperforms. Across retail leadership, healthcare entrepreneurship, investment management, lending, fintech sales, and AI building, my operating pattern has stayed consistent: get close to the outcome, identify the binding constraint, redesign the connected decisions, and make the improvement repeatable.

Top-ranked during tenure across 40+ AEs | 1.5-3x team-average customer value | First 4 enterprise engagements | Built a \$1M+ business

EVIDENCE

I have built and repaired operating systems at materially different scales

At Future Shop, I helped turn underperforming departments into national leaders and helped deliver the first \$2M+ store day across a 130+ store network, with home theatre contributing more than \$1M. I later built Platinum PT, a two-location allied-health company with 20+ staff, from launch to a \$30K monthly run rate within ten months and later \$80K-\$120K per month. Both environments taught the same lesson: leadership cadence, frontline behavior, customer value, capacity, incentives, and unit economics must reinforce one another before growth becomes repeatable.

I use top-line performance to identify the larger system behind it

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 gap showed the opportunity was larger than a better pitch. I reworked price points, packages, larger custom solutions, qualification and disqualification, follow-up, checkpoints, handoffs, support, and account-management requirements. Because the changes proved themselves in customer economics, my influence expanded across the functions that shaped the result and helped lift team-average customer value by roughly 40%.

I know how to make microprocesses operate as one learning system

I also helped develop and operate five live AI applications across prospecting, sales, training, financial-plan guidance, and advisor acquisition. Each prompt, decision tree, configuration change, routing rule, transfer threshold, and scorecard was a microprocess; the advantage came from evaluating where agents and customers succeeded or failed, changing the connected system, and validating the outcome against meetings, handoff conversion, order size, time to close, and revenue. I worked with a 6-8 person product and engineering group on monitoring, releases, and CI/CD coordination while engineering owned production code and infrastructure.

INDEPENDENT PROOF

C2 Lattice, Relay, and Area10 show the same operating method in systems I built independently. C2 makes ownership, authority, dependencies, evidence, and recovery explicit across parallel AI work; its public source is supported by 465 named local tests, separate E2E coverage, 28/28 clean-checkout acceptance, and green hosted CI. Relay does the same for secure private AI workflows, grounding, evaluations, model choice, and cost. Area10 does it for point-in-time data, source health, reasoning, and later decision evaluation. Each turns hidden assumptions into inspectable state and a measurable improvement loop.

WHAT THIS MEANS I help companies move from promising but fragmented traction to a system that compounds: identify the constraint, align the microprocesses, establish ownership and measures, evaluate what is actually happening, and turn each customer outcome into the next product and operating decision.

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