SELECTED IMPACT BRIEF

ROSHAN CHANDRASENA

FINANCIAL SERVICES & AI

I help financial-services companies turn AI capability into trusted workflow, measurable value, and adoption.

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

Financial-services AI has to satisfy four tests at once: economic value, customer trust, compliant workflow, and operational reliability. My background spans SME underwriting and capital allocation, fund formation and trading, financial-planning SaaS, live conversational AI, enterprise commercialization, and point-in-time economic data. That range makes it possible to translate between the decision a customer is trying to improve and the commercial, technical, risk, and adoption conditions required to improve it responsibly.

$4M+ fintech cash collected | Top-five SME lending performer | 39.1M economic observations | Fund founder and head trader

EVIDENCE

I understand financial customers because I have evaluated how businesses use and repay capital

At Thinking Capital/Driven, I reached the top five in a 30+ person commercial organization and originated roughly $1M in funded financing through $25K-$300K facilities across nearly every major industry. Each recommendation required understanding what the capital would unlock, how cash moved through the business, which assumptions made the project viable, and whether repayment matched the use of funds. The result was a broad, practical understanding of business models rather than expertise confined to one product or vertical.

I understand decision risk because I have allocated capital, not only sold financial products

I founded, capitalized, and operated Nota Bene Capital, a Delaware-domiciled investment fund managed from New York. I owned research, portfolio construction, execution, risk controls, and investor materials while coordinating independent operating specialists and a third-party administrator. That accountability sharpened the distinction between a persuasive thesis and a decision that survives uncertainty. Area10 extends the same discipline into software by organizing 39.1M observations across 161,525 economic series around what was knowable at the time—not revised hindsight.

I know how to commercialize financial AI inside real advice and enterprise workflows

At Planswell, I ranked at the top during my tenure across a 40+ AE organization, generated $4M+ in cash collected, and helped guide five live conversational-AI applications. The self-service planning assistant interpreted a 40+ question intake, explained a preliminary plan, answered questions, and moved the appropriate person toward an advisor without replacing licensed judgment. Calls, transcripts, conversation depth, handoff conversion, order size, and time to close became evaluation evidence for improving prompts, decision logic, routing, transfer thresholds, scorecards, and release acceptance. The same approach shaped the company's first four enterprise engagements across financial services.

INDEPENDENT PROOF

C2 Lattice proves I can make agent ownership, authority, dependencies, evidence, and recovery explicit; its public source is supported by 465 named local tests, separate E2E coverage, 28/28 clean-checkout acceptance, and green hosted CI. Relay applies secure, grounded, evaluated, and cost-aware operations to private AI workflows. Area10 preserves point-in-time economic truth and decision lineage. The MIT AI/ML ePortfolio adds externally visible model evaluation and deployment judgment.

WHAT THIS MEANS I help financial-services leaders connect AI to the decision it must improve—then design the customer journey, controls, human handoffs, evaluation loop, and commercial model required for the system to earn trust and produce value.

PROOF LinkedIn | GitHub | AI & ML ePortfolio