

# Kathleen Bartin

AI Transformation · Adoption & Enablement · Value Realization

Millbrook, NY (New York metro) · 314-799-1248 · [kathleen@prodigalparadigm.com](mailto:kathleen@prodigalparadigm.com)  
[linkedin.com/in/kathleen-bartin](https://linkedin.com/in/kathleen-bartin) · [github.com/prodigalparadigm](https://github.com/prodigalparadigm) · [prodigalparadigm.com](https://prodigalparadigm.com)  
Available immediately · New York City · open to hybrid, remote, or relocation

## SUMMARY

---

Enterprise AI transformation, adoption and enablement leader. I build the champions networks and learning infrastructure that make adoption real, ship the systems that prove it, and price what they saved - and because I build the systems myself, skeptical professionals believe me. At a regulated life insurer this summer that meant 77% of the organization engaged - opt-in, with no mandate, on a part-time engagement, against a 50% bar leadership named only afterward. The capability outlived the contract: the client's own senior AI engineer took the enablement library, compiled it into version-controlled skills, and now holds the internal AI build practice. This is the second time I have helped organizations meet a general-purpose technology.

## SELECTED EVIDENCE

---

- **Adoption measured against leadership's own bar:** organizational engagement driven to 77% engaged - opt-in, no mandate, on a sub-20-hour week. No target was set going in; after the engagement closed, leadership put the level they would have been pleased with at 50%. Built through stakeholder discovery and process mapping, role-based training, department cohort activation, office hours, playbooks and weekly executive reporting; group walkthroughs proved the highest-leverage conversion tactic, reaching shift-based and remote staff that individual outreach could not.
- **Capability left behind:** 50+ branded knowledge documents, 15+ training videos, role-based curricula for seven functional teams, seven shared Claude Projects, eight platform skills I packaged plus eight department skills the client's senior AI engineer built from my library, and a curated 56-resource enablement library structured beginner to advanced - a knowledge-transfer infrastructure client staff now own and maintain.
- **Succession, not handoff:** in the engagement's final fortnight the client's senior AI engineer took the enablement library unprompted, converted its prompt guides into version-controlled department skills with compliance guardrails compiled in, credited original authorship, and now holds the internal AI build practice. The engagement did not end with documents; it ended by being built upon.
- **Executive partnership, full cycle:** won, scoped, priced and delivered a CITO-sponsored engagement alone - weekly sponsor reporting on adoption and value metrics, investment framing, and documented conversion of skeptical stakeholders, including the finding that a manager's own non-engagement blocks an entire reporting line.
- **Governance that enabled trust:** advised the Chief Legal Officer on model-level zero-data-retention differences and drove a model's removal from the enterprise picker with the platform-tier upgrade to enforce it; shipped the client's first production AI agent with data classifications, human-in-the-loop controls, logging, responsible-use boundaries and a kill switch.
- **Shipped, not proposed:** a call-center QA scoring system (\$40,000-\$80,000 vendor-equivalent) built with no data-science or engineering team; four agentic builds on Claude Code including an autonomous Slack helpdesk-triage agent on a 30-minute cycle and a synthetic evaluation harness that replaced human testers; and a scheduled automation layer - live leadership adoption dashboard, daily comms triage, weekly tracker reconciliation - integrating ClickUp, Outlook, Teams, Slack, SharePoint and browser automation.
- **Judgment encoded, then measured:** built the Recruiter's Field Kit, an open-source decoder that reads an AI-era job description and reports how many distinct labor markets are inside it, ten-minute screening questions with the real answer and the bluff to listen for, and a verdict graded as received - writing the fix never upgrades the grade. Tuned to my own read and then measured against it: agreement on role counts went from 2 of 6 to 8 of 10. Its central constraint - no automated access to LinkedIn, ever - is enforced by CI on every push rather than asserted in a README. MIT; in testing with practicing recruiters.
- **Value documented:** \$194,500-\$387,500 in market-priced systems, knowledge assets and training on one fixed engagement, plus a recurring-savings estimate of \$156,000-\$421,000 per year - every line traced to a dated sponsor report with stated assumptions and a named owner per layer.
- **Commercial record:** \$200,000+ client revenue from a standing start, 80% close rate, 43% year-over-year transaction growth; separately, 400% of annual reporting quota on a multibillion-dollar capital campaign.

## PROFESSIONAL EXPERIENCE

---

**AI Enablement Consultant - Ceres Life Insurance Company (Contract)** | May 2026 - August 2026

*New York metro · regulated financial services · CITO-sponsored · delivered via Prodigal Paradigm LLC*

- Designed and executed the enterprise AI adoption strategy across seven functional teams: stakeholder discovery, process mapping, opportunity assessment, role-cohort rollout sequencing, and weekly adoption and value reporting to the CITO sponsor. Trained and advised the full C-suite and development leadership.
- Built the learning and enablement program to outlive the engagement - role-based training, office hours, playbooks, cohort activation and champion networks - and transferred it to three named internal owners covering infrastructure and access, library and L&D, and the internal AI build practice.
- Authored the engagement's operating methodology - articulate, test and validate, then automate - which reframed the adoption barrier correctly: the obstacle was rarely technical inability, it was that people had not written down what they actually do.
- Identified, architected and deployed a call-center QA scoring system with no data-science or engineering team - parallel semantic and paralinguistic pipelines, ffmpeg, Whisper and diarization, fully local - and handed it to the operations manager to rebuild herself; migration path to LLM-based judging wired.
- Deployed the client's first production AI agent (Slack helpdesk triage, run observe-first then validated in supervised live runs) with a one-page governance package, plus three further agentic builds on Claude Code across nine weeks.
- Stood up a running automation layer so program telemetry sustained itself: a live leadership adoption dashboard and scheduled agents for daily comms triage and weekly tracker reconciliation, with rate-limit resilience and approval gates.
- Advised the Chief Legal Officer on model-level zero-data-retention differences; drove a model's removal from the enterprise picker and secured the platform-tier upgrade to enforce it. Authored an HR playbook drawing the drafting-versus-screening line under automated-employment-decision rules.
- Brought outsourced legal contract work back in-house with an AI-assisted workflow, valued at displaced outside-counsel spend; diagnosed a multi-day organization-wide AI-platform outage - agentic workloads exhausting shared rate limits and starving interactive use - and designed the remediation architecture separating agentic from human capacity, with approval gates.
- Priced the engagement's value layer by layer from a complete activity log: estimate bands, stated assumptions, and a named owner for each layer.

#### **Principal - Prodigal Paradigm LLC | January 2025 - Present**

*Millbrook, NY · AI adoption consulting and education*

- Champions, not mandates: built "Charity Navigator Lite" from IRS 990 data (94 foundations, \$16B in assets) in under an hour; its technology-skeptical recipient now hosts AI seminars.
- Trained a retail representative with no coding background to ship a working Claude Code application in a single session; unblocked a professional engineer's stalled application in 30 minutes.
- Founded and run "Coffee with Claude" at the Millbrook Public Library - a standing adult AI-literacy community of practice, running since June 2026.

#### **Director of Client Relations (Part-Time) - LaUL Gallery | June 2024 - July 2026**

- Built a \$200,000+ client book from a standing start - no book, no local network, no foot traffic - selling fine jewelry at \$500 to \$25,000+ per transaction; 80% close rate, 50+ high-net-worth accounts; founded the client relations function and recruited and onboarded a successor.

#### **Foundation Services Specialist - The Chamber Foundation | April 2025 - November 2025**

- Designed an eight-lesson youth leadership curriculum from scratch; wrote approved one-, two- and three-day workforce bootcamp curricula; orchestrated 80+ volunteers across four committees.

#### **Prospect Researcher to Senior Researcher - Washington University in St. Louis | 2009 - 2012**

- Exceeded annual reporting quota by 400% on a multibillion-dollar capital campaign; pioneered international prospecting and social-media research methods that later became industry standard.

Family sabbatical, 2012-2024 - full-time parent; intermittent contract roles. Earlier: Technology Trainer and Sales, New Horizons / Productivity Point International - sold and delivered enterprise computing training during the first wave of corporate adoption.

### **FRAMEWORKS AND METHODS I BRING**

---

- **The Underwriting Principle:** you would not price a product on unvalidated assumptions; do not automate a workflow on unarticulated knowledge. Three phases - articulate, test and validate, automate.
- **Adoption maturity ladder:** Equipped (tools installed, connectors and add-ins live), Oriented (knows the surfaces, the library, the settings, what a skill is), Builder (has shipped one skill end to end), Claude Code ready (graduated from skills to agentic builds). Used to define "engaged" operationally so every monthly reading measures the same thing.
- **The Rule of Three:** if something has been explained three times, it belongs in a shared Project - the mechanism that turns one person's know-how into durable organizational capability.
- **Make the client the hero:** every asset built to be usable without the consultant present, and every engagement designed to end in succession rather than dependency.

## PORTFOLIO - BUILDS AND RESEARCH

---

- **Twelve public repositories**, MIT-licensed, across two GitHub accounts; the four suites below were built after the Ceres engagement closed ([github.com/prodigalparadigm](https://github.com/prodigalparadigm)):
  - **servicenow-mcp** - an MCP server over ServiceNow ITSM with field allowlists pushed down as sysparm\_fields so unvetted columns are never transferred, a hard per-call record cap, token-bucket rate limiting with 429-aware exponential backoff honoring a clamped Retry-After, read-only by default enforced twice (absent from the tool registry and refused at the service layer), and correlation-id idempotency so a replayed create does not file the incident twice.
  - **stagegate** - a dependency-free Python library that runs each agent capability at a declared stage - observe, suggest, act - resolved from a reviewed TOML policy rather than code, with one hash-chained append-only audit record per call, redaction applied before any argument reaches a sink or a human, fail-closed approval timeouts, and a kill switch that degrades to shadow mode instead of raising.
  - **promptproof** - generates synthetic conversations across six axes, including scripted multi-turn-drift ladders that install a premise over two or three turns before cashing it in, runs them against multiple models, scores with a file-based LLM-judge rubric, and reports per model which specific instruction broke and on which turn, with hard-versus-soft columns and errors never folded into failures.
  - **callscope** - a fully local two-speaker call-QA pipeline: ffmpeg normalization to 16 kHz mono, energy VAD with hysteresis, MFCC-plus-log-F0 embeddings, duration-weighted 2-means initialized at the PC1 weighted median, a silhouette guard that collapses to one speaker when the split is not supported, and semantic and paralinguistic tracks scored independently so a disagreement between them is informative.
  - Every suite runs offline with no network and no credentials, and CI runs pytest, ruff and strict mypy on every push.
- **Teaching, packaged and maintained:** a 30-minute hands-on lab on Claude tool use for GTM teams - the request, tool\_use, tool\_result loop - taught against a policy-lookup scenario that answers the first question every enterprise buyer asks: will it make things up about our data? Ships with a pre-class smoke test that must print four passes before anyone teaches it, captured transcripts so the room keeps working through an API outage, instructor notes naming the three likely failures, and a release-cycle maintenance checklist. MIT.
- **Production personal stack:** nine Model Context Protocol (MCP) servers and a hybrid-retrieval RAG system (dense plus sparse, reciprocal-rank fusion, 64,000+ document chunks), operated daily in my own practice.
- **RAVEN:** open-source Claude-powered agent that learns aesthetic taste from a curated corpus and hunts dealer inventory; 17 signal layers; MIT license; 53+ unique GitHub clones.
- **Longitudinal AI-behavior research:** pre-registered measures over a 495-conversation, 20,000-turn corpus; Kaggle / Google DeepMind Metacognition Track (2026) submission.
- **Safety differential finding:** AI security hackathon, New York City 2026 - identical surveillance-agent prompt across two frontier models: model choice is a safety decision, not a branding decision.

## SKILLS

---

- **Adoption and change:** AI adoption strategy, adoption measurement and value realization, change management, communities of practice, champions networks, train-the-trainer, journey mapping and friction analysis, role-cohort rollout design, stakeholder discovery, process mapping, opportunity assessment.
- **Enablement and learning:** instructional design, technical curriculum development, training delivery, hands-on labs and workshops, office hours, AI literacy programs, playbook and SOP development, technical translation for non-technical audiences, executive briefing.
- **Governance and advisory:** C-suite AI strategy, investment framing, executive sponsorship and stakeholder buy-in, platform and console governance, data classification, human-in-the-loop design, responsible-use policy, compliance-aware deployment in regulated industries, model selection and retention posture (enterprise-picker governance, retention-posture comparison, cross-model safety differential), tool and connector evaluation, vendor and partner evaluation.
- **Technical:** Claude (Enterprise, apps, Claude Code, API), ChatGPT, Gemini, multi-model evaluation and governance, prompt engineering, agent design, MCP server development, RAG (Qdrant, ChromaDB), Python, Git and GitHub, ffmpeg, Whisper, speaker diarization, agentic browser automation (Claude in Chrome), SharePoint, Microsoft 365, Slack, Teams, ClickUp.

## EDUCATION, CREDENTIALS AND COMMUNITY

---

- BA, Philosophy - Excelsior University. Coursework, University of Chicago. Graduate coursework in Actuarial Mathematics, University of Missouri-Columbia.
- Eleven Anthropic Academy certificates (2025-2026). Google AI Essentials (2025).
- US Navy, Honorable Discharge. Licensed EMT (New York). Board member, Millbrook Business Association. Rotary Exchange Counselor. Millbrook Hook and Ladder.