BREAKVECTOR LABS

TRANSFORMATION CONSULTING

FOR BANKS, CARD ISSUERS, INSURERS AND ASSET MANAGERS

Governed AI transformation, from pilot to production.

A data-first delivery method, governance built into every gate, and change management that makes adoption stick, run by one accountable program lead.

STAGE 0 Mobilize and govern	STAGE 1 Frame the use case	STAGE 2 Prove the data	STAGE 3 Prepare the data
STAGE 4 Build	STAGE 5 Validate	STAGE 6 Deploy and operate	STAGE 7 Scale and sustain

Eight stages, seven governance gates, three workstreams running throughout.

BUILT ON FOUR RECOGNIZED BODIES OF PRACTICE

PMI-CPMAI
Data-first AI delivery lifecycle

IAPP AIGP
AI governance, law and risk

Prosci ADKAR
People-side change and adoption

PMP
Program leadership and value

Jiby Manuel, PMP · Prosci Certified

[Email] · [Phone] · [Website]

THE CHALLENGE

AI rarely stalls because the model is weak. It stalls everywhere around the model.

1 **No owner, no metric, no tier**
Use cases start without a business owner or risk tier, then stall at the first risk review.

2 **Data that won't pass review**
Data isn't lawful, documented or representative enough for validation or fair-lending tests.

3 **Governance arrives last**
Controls bolted on at sign-off add months of rework instead of shaping the design.

4 **Front lines don't adopt**
Fraud, Risk and Operations teams don't trust the tool, so the business case never lands.

5 **The rules keep moving**
2026 alone rewrote US model risk guidance, moved EU AI Act dates and replaced Colorado's AI law.

= **The result**
Dozens of pilots, a handful in production, and examiners asking for evidence nobody kept.

OUR APPROACH

Delivery, governance and adoption as one program

TYPICAL APPROACH	THE BREAKVECTOR LABS APPROACH
Pick a tool, then look for a use	Pick a use, prove the data, then choose the tool
Governance as a final sign-off	A governance gate at the end of every stage
Same controls for every use case	Risk tiers: fast lane for low risk, full rigor for credit and fraud
One big launch	Small, complete iterations with an honest go/no-go
Training at go-live	Adoption plan from week one, measured with ADKAR
Separate firms for strategy, build, risk, change	One accountable program lead across all four

WHO WE SERVE
Banks, card issuers, payments, insurers, wealth and asset managers

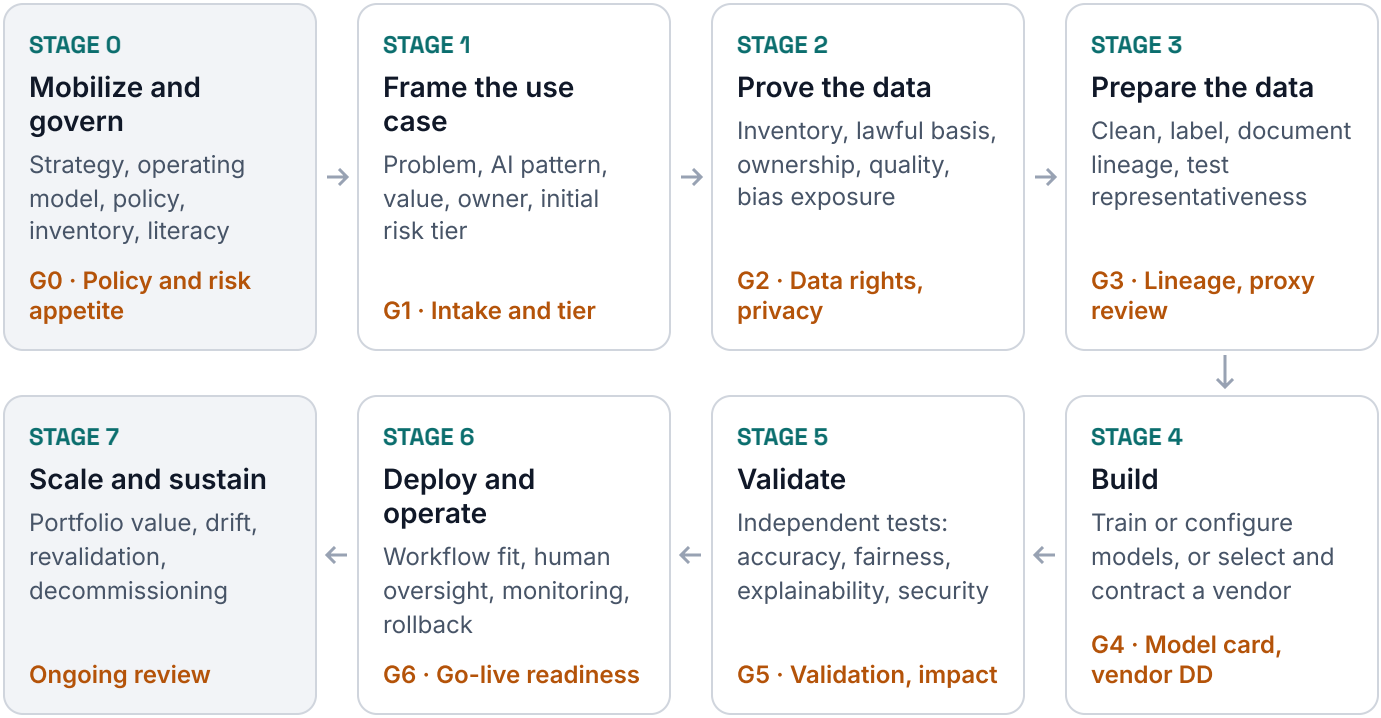
WHO WE WORK WITH
Chief Risk, Fraud, Operations, Data and Transformation officers

WHERE WE FOCUS
Risk, Fraud and Operations, where AI meets customers and examiners

THE BREAKVECTOR LABS FRAMEWORK

The Governed AI Lifecycle

PMI-CPMAI's six data-first phases form the spine. IAPP AIGP controls sit on every gate, so each go/no-go is also a governance decision. Stage 0 runs once per firm; stages 1 to 6 repeat for every use case.



THREE WORKSTREAMS RUN THROUGH EVERY STAGE

Governance and risk · AIGP	Inventory, risk tiering, policies, impact assessments, third-party oversight, monitoring, incidents
Change and adoption · ADKAR	Impact analysis, sponsor roadmap, change-agent network, role-based training, reinforcement
Program and value · PMP	Portfolio intake, prioritization, budget, benefits tracking, RAID, executive reporting

RISK TIERING SETS THE WEIGHT OF CONTROL

Tier 1 · High

Credit decisions, transaction-blocking fraud, AML scoring

Full independent validation, fairness and adverse-action tests, quarterly monitoring

Tier 2 · Medium

Analyst decision support, dispute triage, call summaries

Targeted validation, documented testing, monthly monitoring

Tier 3 · Low

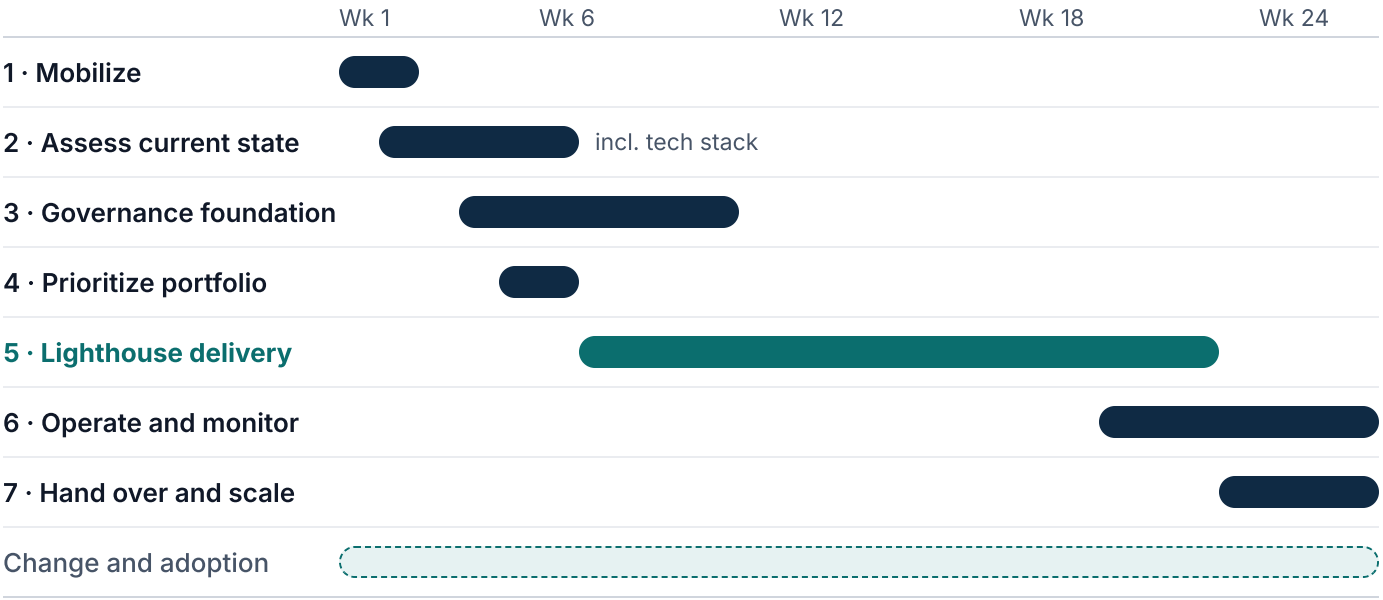
Internal productivity genAI, no customer decision

Acceptable-use policy, data-loss controls, usage monitoring

END-TO-END PLAYBOOK

Seven plays, 26 weeks, lighthouses in production

The first wave assesses, builds the governance foundation and takes two or three lighthouse use cases to production. Durations are planning defaults, tuned to each client.



INSIDE PLAY 5: EACH LIGHTHOUSE, SPRINT BY SPRINT

WK 7–8	WK 9–10	WK 11–12	WK 13–16	WK 17–19	WK 20–22
Frame	**Prove data**	**Prepare**	**Build**	**Validate**	**Deploy**
KPI baseline, process map, tier	Profiling, lawful basis, bias scan	Pipelines, lineage, proxy review	Baseline, model card, guardrails	Fairness, red-team, UAT	Pilot group, monitoring, training

By day 90

· Maturity scorecard and stack heat map
· AI governance committee chartered and meeting
· AI policy, genAI standard and risk tiering
· One AI inventory, vendor AI included
· Ranked backlog; lighthouses in build

By week 26

· Lighthouses live with an evidence pack each
· Monitoring, incident playbook, revalidation triggers
· Benefits measured against baseline
· AI center of excellence designed
· Wave-2 roadmap approved

TECHNOLOGY STACK ASSESSMENT

Can you build, run and evidence AI safely at scale?

A 3 to 4 week review that scores eight layers from 1 to 5 against a target set by your ambition and risk mix, then shows which gaps block which use cases.

1 Data foundation
Cataloged data with lineage, quality scores, owners

2 Data access and privacy
Purpose-based access, tokenized card and PII data

3 ML and AI platform
Standard environments, feature store, model registry

4 Generative AI layer
LLM gateway, governed retrieval, guardrails, evals

5 MLOps and LLMOps
Versioned release path with approvals and rollback

6 Monitoring and observability
Drift, fairness, cost and safety alerts on thresholds

7 Security and resilience
Prompt injection, leakage, poisoning; red-teaming

8 Governance and risk tooling
One AI inventory; intake-to-evidence workflow

Also in scope: integration with card processing, case management, CRM and contact center, and AI embedded in existing SaaS.

SCORING RUBRIC

1 · Ad hoc	2 · Emerging	3 · Defined	4 · Managed	5 · Optimized
Manual, no standards	Some shared tools	Standard, evidence on request	Automated, evidence by default	Policy enforced in code

Target: 3 or better on every layer before a Tier 1 use case goes live; 4 on layers 5, 6 and 8 to scale.

QUESTIONS WE ASK

Can you trace a live model's training data back to source?

Which LLMs are approved, and through what gateway?

Can you redeploy last quarter's model version in a day?

What alert fires if fraud approval rates shift 5 points overnight?

RED FLAGS

No inventory, or one that skips vendor AI

Live models with no owner or monitoring

Customer data in public genAI tools

Builders validating their own models

You receive: Heat map, current vs target · Gaps ranked by risk and blocked use cases · Build, buy or reuse options · Target architecture and roadmap

REGULATORY READINESS

2026 changed the rules. Every gate maps to them.

Each control in the lifecycle is tagged to the rule it satisfies, so you can show an examiner or auditor exactly where a requirement is met.

APRIL 2026 · US BANKING

New model risk guidance replaces SR 11-7

Fed SR 26-2, OCC 2026-13, FDIC FIL-15-2026. Scaled to risk; more on vendor models. GenAI and agentic AI left for separate guidance.

DEC 2027 · EU AI ACT

High-risk duties, incl. credit scoring, deferred

Digital Omnibus moved Annex III from Aug 2026. Chatbot and synthetic-content transparency applies from Aug 2026.

JAN 2027 · COLORADO

SB 26-189 replaces the Colorado AI Act

Notice before consequential automated decisions; explanation, correction and human review within 30 days of an adverse outcome.

RULE-TO-GATE MAP

RULE OR FRAMEWORK	WHAT IT ASKS OF AI	WHERE
Model risk guidance (2026)	Sound development, independent validation, monitoring, governance	G4–G5, S7
ECOA / Reg B, FCRA	Specific adverse-action reasons; no disparate impact; data accuracy	G2, G3, G5
UDAAP	Chatbots, collections and marketing AI must not mislead	G5, G6
Third-party risk guidance	Due diligence, contract rights and oversight for vendor AI	G4
GLBA, PCI DSS	Protect customer and card data in pipelines and prompts	G2, stack
EU AI Act	Literacy, transparency, high-risk system obligations by role	S0, G1, G5
NIST AI RMF + GenAI profile	Govern, Map, Measure, Manage	S0–S7
ISO/IEC 42001, 42005	Certifiable AI management system; impact assessment	S0, G5



NIST AI RMF functions mapped to lifecycle stages.

Status as of November 2026. A briefing aid, not legal advice; confirm specifics with counsel.

HOW WE ENGAGE

Start small. Expand on proof.

Five modular offers. Most clients begin with the assessment, then add a governance foundation and a lighthouse wave.

AI Readiness and Stack Assessment 4–6 weeks · start here	Six-dimension maturity scorecard, technology heat map, regulatory applicability map, ranked use-case backlog, 18-month roadmap.
AI Governance Foundation 8–12 weeks	Committee charter, AI and genAI policy, risk tiering, AI inventory, impact assessment, RACI, literacy program; NIST and ISO 42001 aligned.
Lighthouse Delivery 12–16 weeks per wave	Two or three use cases taken through every gate to production, with change management and benefits tracking.
AI Transformation Office 6–18 months	Fractional or full-time program leadership: PMO, value tracking, governance operations, change, executive reporting.
Adoption Accelerator 6–10 weeks	For tools that are live but under-used: Prosci impact assessment, sponsor roadmap, change agents, role-based training, adoption metrics.

Fixed scope Assessments, foundations	**Retainer** Days per month	**Interim leader** Embedded 6–12 months	**Advisory seat** Monthly with committee

YOUR PROGRAM LEAD

Jiby Manuel

PMP · Prosci Certified Change Practitioner

· Program leader in banking and card: Risk, Fraud and Operations transformation
· Led an AI program in client reporting at Nuveen
· Took an AI use case into production for a cloud warehouse platform
· Built PriorVue AI, a portfolio layer on Jira Cloud

Book a 45-minute AI readiness conversation

[Email]
[Phone]
[Website / LinkedIn]
Charlotte, NC