

Executive guide to forging AI value

A breakpoint-first framework for building AI strategy that delivers real business outcomes





break·point / brāk·point /

Where decisions stall. Where information hides. Where value leaks.

The AI pressure is real — and right now, it’s nearly impossible to separate signal from noise.

When a path forward does emerge, new questions surface fast: Who trusts the model’s output? Who’s accountable when it’s wrong? How does work get redesigned around it? There’s no established playbook. That’s why more than 80% of AI projects fail to reach production — twice the rate of traditional IT projects.

After working with enterprises on this every day, we know it’s rarely the technology that breaks. What fails is the connection between AI capabilities and actual business value.

The organizations succeeding with AI don’t ask “What can AI do?” They start by identifying their breakpoints — the specific places where processes break down, friction accumulates, and value escapes. Then they put AI and agentic technologies to work fixing them.



Find your breakpoints: Three gaps to success

The journey from identifying a breakpoint to fixing it at scale requires closing three gaps. Most AI initiatives stall because they underestimate at least one.



Vision Gap

The problem: The organization lacks clarity on what success looks like. Teams identify dozens of breakpoints but can't prioritize. Committees spend months defining governance while competitors ship products.


The test: Can your leadership team think ahead 18 months and articulate a one-sentence headline about your AI success? If not, you have a **vision gap**.



Trust Gap

The problem: The organization doesn't trust data to fuel AI. It's incomplete, inconsistent, siloed, or ungoverned. Leaders assume their data is "clean enough" without validation — then wonder why models produce unreliable results.

The test: Do you know who owns your critical data assets? Who's accountable for quality? Who can approve its use? If you can't answer confidently, you have a **trust gap**.



Scale Gap

The problem: Infrastructure can't run AI in production. Pilot environments are small, clean, and forgiving. Production is massive, messy, and unforgiving — with real security requirements and real cost scrutiny.

The test: How long would it take to deploy a new model to production today with full security compliance? If the answer is "months" or "we're not sure," you have a **scale gap**.

Start with your headline: Building an AI strategy that delivers

AI has been a board-level priority since 2023, yet most organizations have yet to achieve real traction. The root cause is almost always strategic.



It's 18 months from now.

Your industry's most respected publication wants to interview you about your organization's AI journey. They're going to write a feature highlighting your success.

What's the headline?

What have you accomplished? Why do your customers care? What metrics have you dramatically improved?

The answer can't be "We implemented AI across the organization" or "We're using AI to be more efficient." You need a headline that's specific, measurable, and customer centered — like these:

- **Regional Bank Cuts Loan Approval Time from Five Days to Four Hours**
- **Manufacturer Reduces Customer Onboarding by 60%, Boosts First-Year Retention 25%**
- **Healthcare System Eliminates 40 Hours/Week of Clinician Documentation Burden**



If your leadership team is still struggling to identify a compelling North Star, you've identified the real barrier to AI ROI. That's the vision gap — and closing it is the first step.

Five questions to find your "why"

AI success starts with a conversation about the business, not technology. As you work to surface your organization's breakpoints, start here:



If we could accomplish only one thing with AI in the next 18 months, what would move the needle most for our customers?



Where are our highest-value employees spending time on tasks that don't require their expertise?



What decisions do we make repeatedly that would benefit from faster, better information?



Which of our current processes creates the most friction for customers or employees?



Do we have executive alignment on what AI success looks like — not in theory, but in measurable outcomes?

Fostering trust: Your AI is only as good as the data behind it

The percentage of companies abandoning the majority of their AI initiatives before they reach production surged from 17% in 2024 to 42% in 2025. Our experience suggests that such failures are often rooted in the invisible infrastructure around AI: data governance, quality controls, ownership structures, and integration with actual business workflows.

Where AI really breaks down

RAND Corporation interviewed 65 data scientists and engineers to find the root causes of AI project failure:

#1 Misunderstanding what problem needs solving

Technical teams don't understand the business context.
Business stakeholders can't articulate the actual problem.

#2 Organizations don't have the data they think they have

Data is incomplete, inconsistent, siloed, or stale.
Organizations assume it's "clean enough" without validation.



Five roles that must be defined

Before any AI initiative moves forward, you need clear answers to several questions: Who owns this data? Who's accountable for its quality? Who can approve its use? Answering those questions requires defining and empowering five roles:



Chief data officer. Accountable for overall governance program success. Sets strategy, coordinates across teams, and secures executive alignment.



Data owner. Accountable for domain data quality, compliance, and proper use. Defines access policies and approves use cases.



Data steward. Responsible for day-to-day governance compliance. Monitors integrity, enforces quality rules, and bridges business and IT.



Data custodian. Responsible for technical guardrails and system reliability. Implements security, storage, backup, and access controls.



AI governance lead. Authority to block AI deployments that fail governance criteria. Owns model documentation, bias testing, and explainability standards.

Building the foundation in 90 days

Data governance isn't a one-time project; it's an operational capability. But if you focus on the right priorities, you can establish the critical foundation in 90 days:

Days 1–30: Establish the charter


Define AI vision with measurable outcomes • Conduct readiness assessment • Inventory data assets • Establish governance charter with executive sponsorship • Designate owners, stewards, and custodians

Days 31–60: Build the controls

Certify priority datasets (~10) • Implement data lineage • Deploy role-based access controls and audit logging • Launch "golden path" templates • Create role-specific training

Days 61–90: Prove it works

Pilot AI on low-risk use case with full governance • Demonstrate audit pack generation in <5 minutes • Establish model registry • Define bias testing protocols • Set up drift detection



From pilot to production: Infrastructure that scales with AI

You've defined your AI strategy. You've addressed your data governance. Your pilot is working beautifully in the lab. Now comes the hard part.

The gap between a successful proof of concept and a production system running at enterprise scale is where AI initiatives often falter. The most common cause: underestimating operational complexity.

Pilot environment		Production reality
Small, clean data	→	Massive, messy, sensitive data
Single use case	→	Competing workloads
Dedicated resources	→	Shared infrastructure
Relaxed security	→	Enterprise security requirements
Cost is secondary	→	Cost under scrutiny

Three barriers to scale

Organizations that successfully scale AI have learned to address three interconnected challenges that can sink production deployments:

Complexity

AI infrastructure isn't a single system. It's dozens, including model training, deployment pipelines, monitoring, retraining, and version control. Each brings different security architectures, different scaling characteristics, and different failure modes. Integrating them is where projects can stall.

Security and governance

Production AI touches sensitive data at scale. It requires role-based access across multiple tools, each with its own security model. Compliance requirements multiply. And AI-powered attacks are getting more sophisticated, targeting the infrastructure itself.

Unpredictable costs

GPUs are expensive and often idle. Cloud costs scale unpredictably. Inference costs at production volume look nothing like pilot costs. Without discipline around resource allocation, AI initiatives blow through budgets before they deliver value.

The infrastructure decision tree

Before you can run AI at scale, you need clear answers to four strategic questions:

Where will AI run?

Will you opt for cloud-first, on-premises, or hybrid? The answer should be driven by data residency requirements, latency needs, and regulatory constraints, not preference.

What's your GPU strategy?

Do you have the capital for owned compute? The workload predictability for reserved capacity? Or does elastic cloud scaling better match your usage patterns?

How will you handle the speed of life?

What systems will your AI agents need to interface with, and where do those systems live? Real-time requirements change everything.

What's your security integration plan?


How will you maintain consistent access control, audit logging, and compliance across the full ML lifecycle? And how will you defend against AI-powered attacks?

The path forward

OnX Forge AI implementation roadmap

Strategy, data, and infrastructure don't operate in isolation. The organizations that forge AI value treat them as a coordinated lifecycle — moving from vision to secure, enterprise-scale operations.

OnX Forge AI supports your progression from strategic foundation to AI that's embedded in your enterprise DNA. Each tier builds on the last, and each delivers tangible outcomes before the next begins.


Tier 1 

Advisory (visioning and workshops)

Purpose: Establish a strategic foundation and align AI initiatives to business outcomes.

Offerings: AI Consulting

Deliverables: AI roadmap • ROI framework
• Prioritized use case pipeline

Tier 2 

Data & infrastructure readiness

Purpose: Build the foundation for AI by ensuring data and infrastructure are enterprise-ready.

Offerings: Data Readiness, AI Infrastructure

Deliverables: AI-ready data estate + scalable, secure infrastructure platform


Tier 3 

Deployment

Purpose: Deliver tangible business value through real AI use cases.

Offerings: GenAI Development, AI Jumpstart, Vertical Solutions

Deliverables: Working pilots → production deployments with expansion roadmap

Tier 4 

Secure scale

Purpose: Operate safely and securely at enterprise scale.

Offerings: AI-Native Security, Governance and Compliance

Deliverables: AI becomes part of enterprise DNA — secure, governed, continuously improving

Find your breakpoints. Forge real value.



The first step is identifying where AI can create real impact for your business — not chasing technology for its own sake. OnX Forge AI Strategy Sessions help leadership teams find the breakpoints that matter, prioritize ruthlessly, and build a roadmap that delivers.

Start a breakpoint-first conversation.

Contact OnX

About the Expert



Celio Casadei

Senior Vice President, AI

Celio leads OnX's AI practice across Canada, bringing decades of digital transformation experience to the work of turning enterprise ambition into production-ready AI outcomes. He partners with enterprise clients to design and scale programs that span strategy, cloud, data, and generative AI — with a clear focus on measurable business value, responsible governance, and long-term operability. Organizations engage Celio's team for end-to-end accountability, from executive alignment through to ongoing optimization.

With a career built at the intersection of technology and business transformation, Celio brings deep expertise across banking, insurance, telecommunications, and the public sector. He has led large-scale digital transformation initiatives both as an executive and as a strategic advisor, specializing in helping organizations extract maximum value from their data assets — including pioneering work in data integration, ERP modernization, and enterprise AI adoption. His work consistently bridges the gap between technology capability and business strategy, delivering bottom-line impact in complex, high-stakes environments.

Celio is a recognized advocate for responsible innovation — championing emerging technologies, inclusive leadership, and talent development as interconnected pillars of sustainable organizational change.

Sources

RAND, The Root Causes of Failure for Artificial Intelligence Projects and How They Can Succeed: Avoiding the Anti-Patterns of AI
S&P Global Market Intelligence, AI experiences rapid adoption, but with mixed outcomes – Highlights from VotE: AI & Machine Learning | S&P Global