

AI Strategy Leadership Clinics - A Decision-Ready Playbook

# From AI Ambition to Production- Grade Strategy

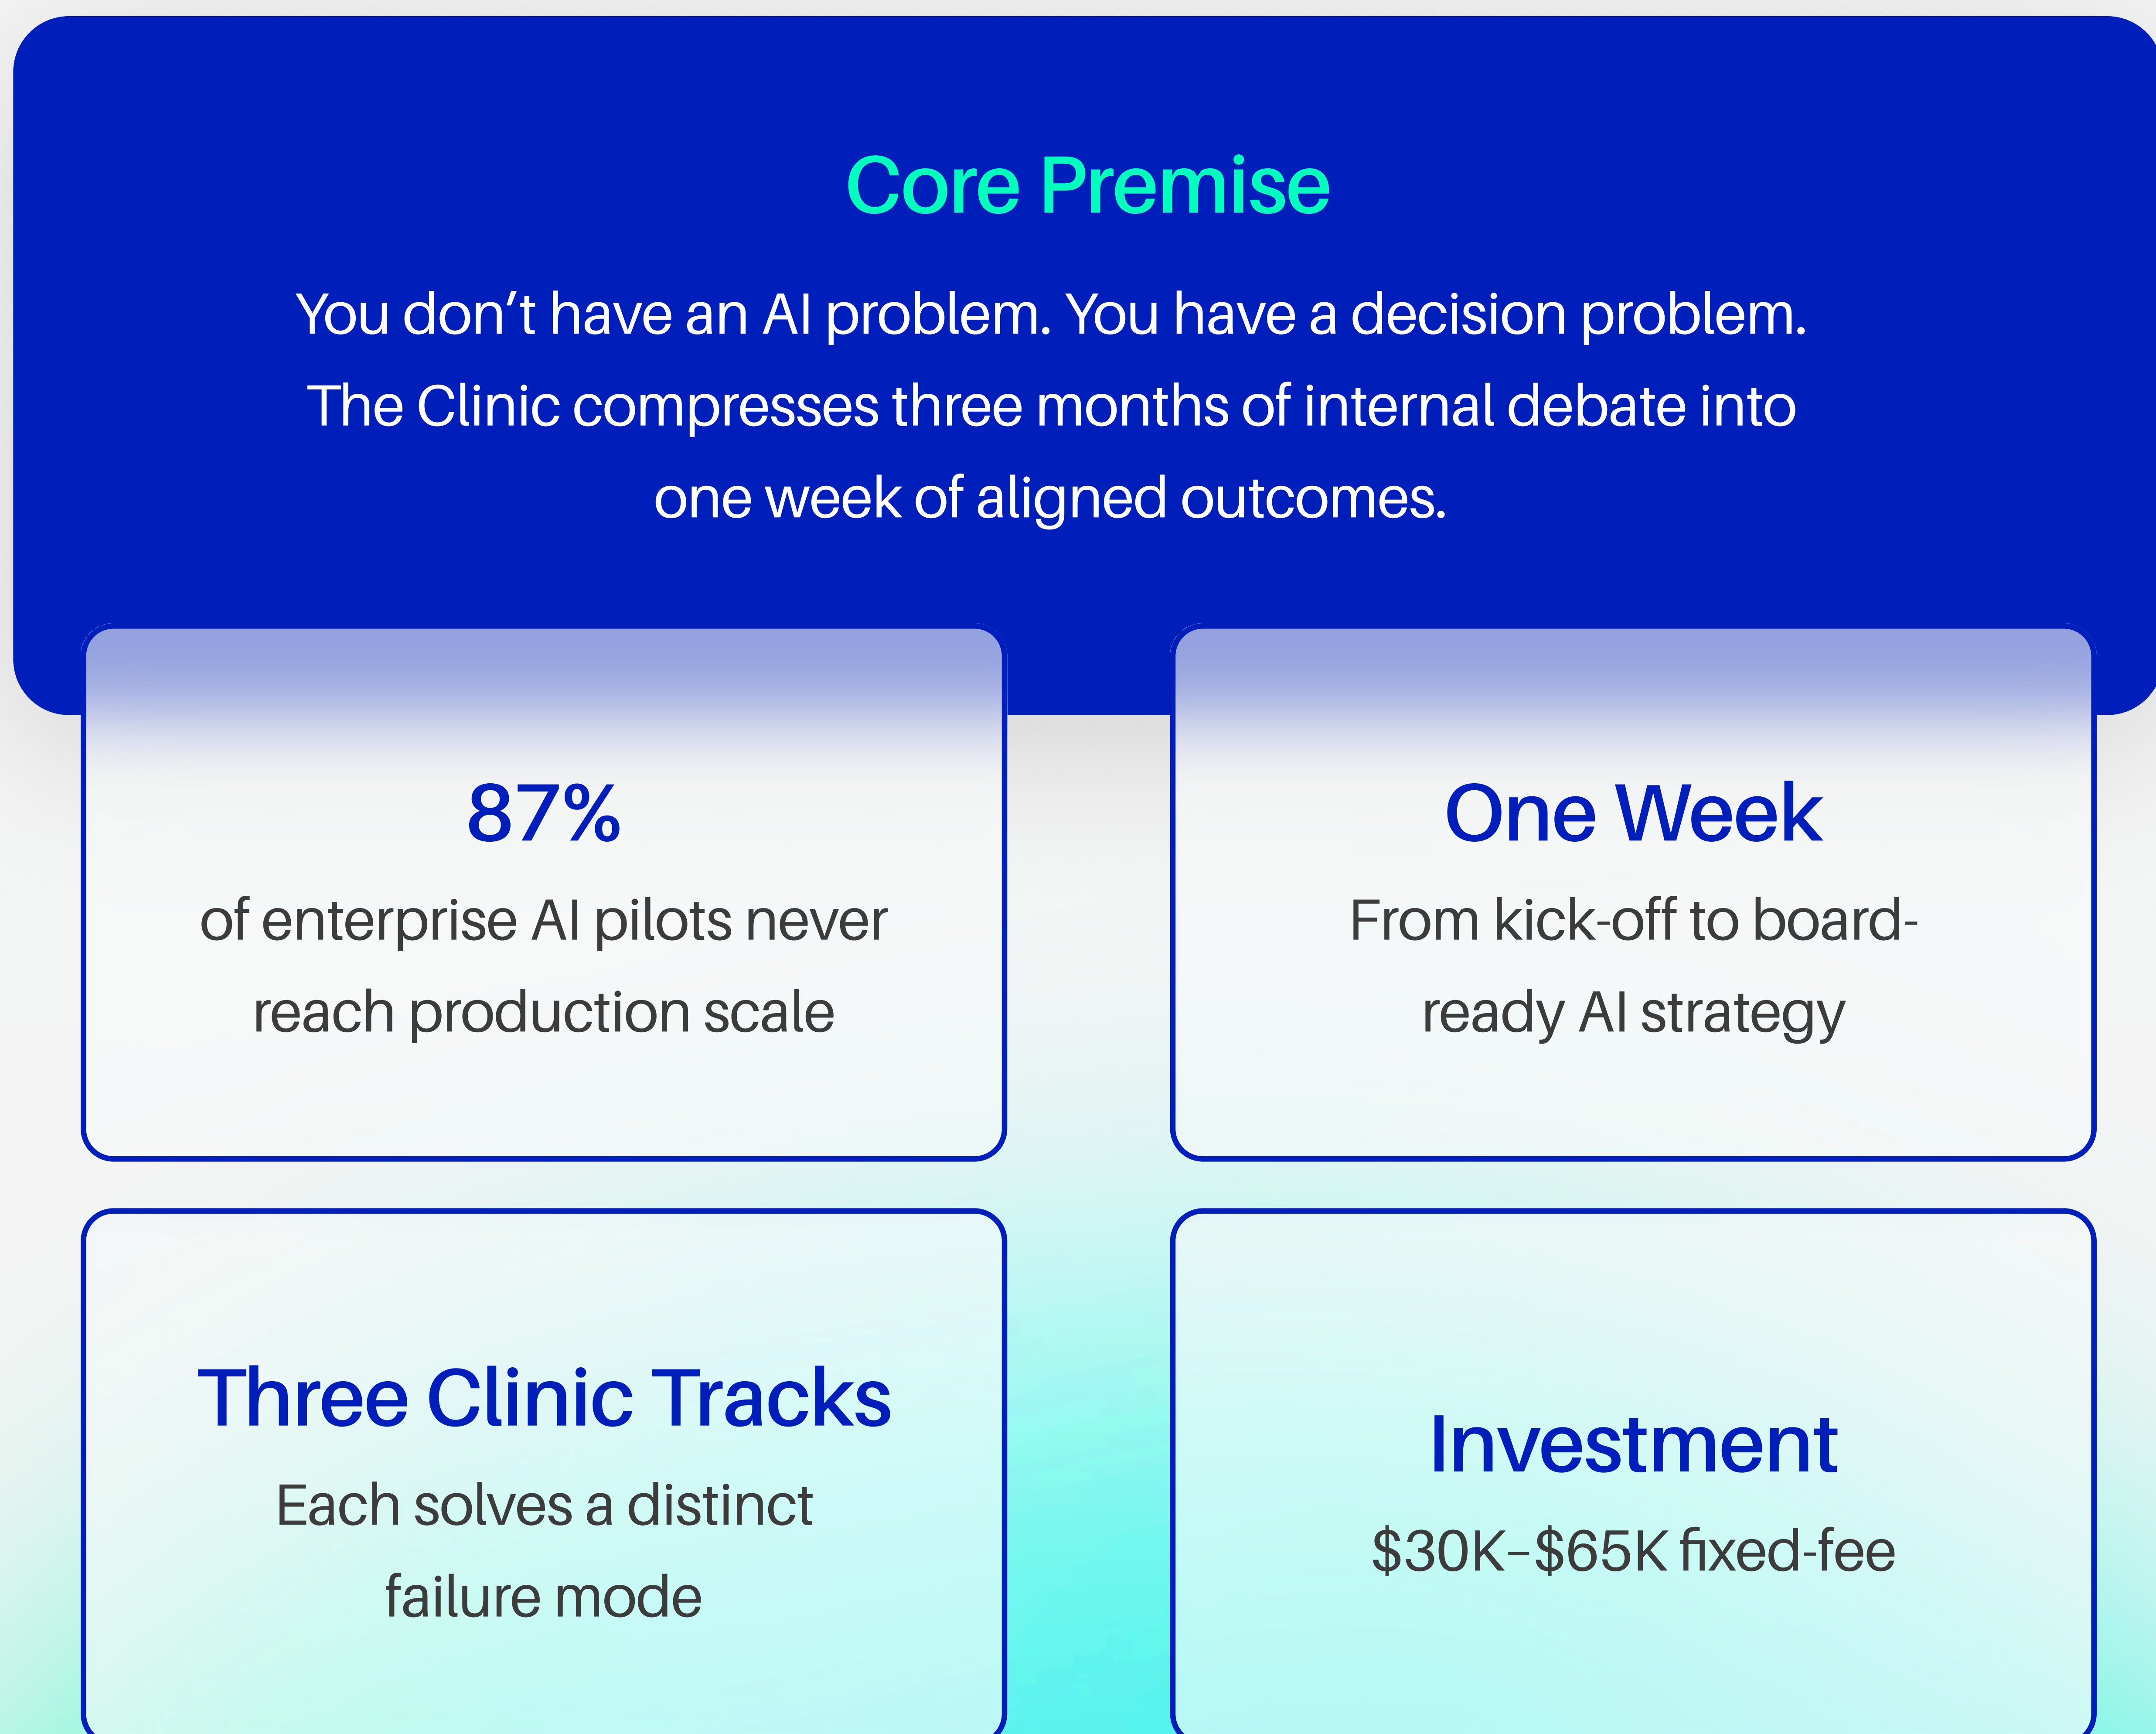
A practitioner's guide to escaping the PoC Trap, closing the strategy-to-execution gap, and generating measurable EBIT from AI.




## The Situation in One Page

There's a myth that most enterprises are failing to start AI programs. In reality, they are failing to finish them. The PoC loop is a hard-to-swallow reality: teams prove technical feasibility in isolation, but stall at the boundary between experimentation and production. The reasons are predictable, but they are also fixable.

The shift to Agentic AI - autonomous, multi-step systems that operate across your data and processes - demands more than new capabilities. It requires an aligned platform strategy, a single governing thesis, and the discipline to kill underperforming experiments before they consume the budget and leadership attention that should be going to production.



## What You Leave With



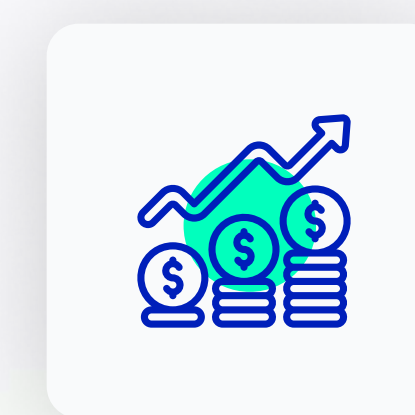
### Strategic

- Validated \$1B question\*
- Competitive Moat Analysis
- Board-ready Executive Briefing Pack



### Technical

- Agentic Architecture Blueprint
- Governance Model
- Integration Roadmap



### Financial

- EBIT Impact Models
- Prioritized Use Case Funnels
- Board-Ready Business Cases

**Note:** The '\$1B Question' is not a revenue target - it's a structured exercise that compels leadership to name, specifically, where AI will create or destroy the most value in the business over the next few years.

## The Problem – Why AI Programs Stall

### Three Failure Patterns. One Root Cause.

Before choosing a path forward, leadership needs an honest account of where the AI program is breaking down. Across hundreds of enterprise engagements, the same three patterns recur frequently and at the same time.



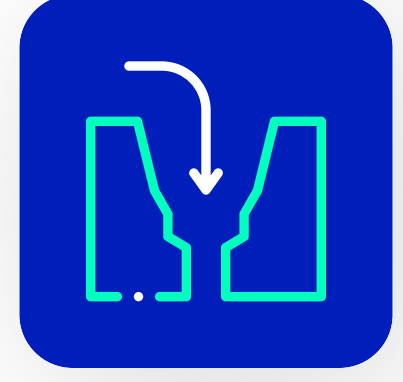
### The “Tool” Trap – Lack of Strategic Vision

#### What it looks like

Pilots are running, vendors are engaged, and teams are actively experimenting. But the underlying strategic question remains open: where, specifically, will AI create or protect value at scale? When that stays unanswered, initiatives multiply rather than compound - each defensible on its own, collectively pulling in different directions.

#### The signal to watch

Multiple AI initiatives are live, each reporting progress against its own metrics. But there's no single view of what the portfolio is delivering collectively, and no shared language for deciding which initiatives deserve more investment.



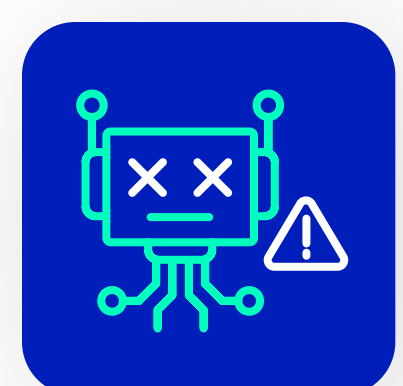
## The “Silo” Trap – Architectural Gaps

### What it looks like

Different teams build their own AI tools to solve their own problems — and for a while, that feels like progress. But those solutions don't talk to each other. What began as moving quickly quietly becomes a drag: teams spend more effort maintaining what exists than building what comes next.

### The signal to watch

The teams doing the most with AI are spending a growing share of time on integration and maintenance rather than new capabilities. Each addition requires more custom work than the last.



## The “Feasibility” Trap – Risk Blindness

### What it looks like

The use case is strong, the business case is solid, and the team is motivated. Then, a few months in, questions surface around data quality, regulatory fit, or adoption. Not unusual questions - just arriving later than comfortable, when changing course carries a real cost.

### The signal to watch

Legal and compliance are consulted after architecture decisions are made. Data readiness is assumed rather than validated. The work gets done - but slower and more expensively than planned.

# Where Does Your Organization Stand?

Mark the option that most accurately reflects your current state.

Dimension	Not Yet (1)	Emerging (2)	Operational (3)	Score
<b>Leadership Alignment</b> (CxO team shares a unified AI ambition)	No shared vision, each exec defines AI success differently	Directional alignment, some shared language, no formal strategy	Board-approved AI thesis with defined owners and metrics	<input type="checkbox"/>
<b>PoC-to-Production Track Record</b> (Ability to scale AI beyond pilots)	Pilots stuck, no reliable path to production	Some production wins, scaling inconsistent	Repeatable process to move idea to production at speed	<input type="checkbox"/>
<b>Platform &amp; Architecture Clarity</b> (Defined AI platform and orchestration model)	No unified architecture	Emerging standards in one function, no enterprise coherence	Unified agentic platform with documented orchestration	<input type="checkbox"/>
<b>Governance &amp; Risk Posture</b> (Trust-by-Design and compliance protocols)	No formal AI governance, compliance is reactive	Policies being drafted, legal involved in some decisions	Risk architecture embedded from the design stage, audit-ready	<input type="checkbox"/>
<b>Data &amp; Integration Readiness</b> (Data quality, accessibility, and labeling)	Data siloed, inconsistent, pipelines built per-project	Core datasets accessible, gaps in quality and real-time availability	Clean, federated data with MLOps pipelines	<input type="checkbox"/>
<b>Use Case Prioritization</b> (Quantified, validated pipeline of opportunities)	Long list of ideas, no structured scoring or financial model	Top ideas informally ranked, EBIT impact estimated, not modeled	PRIME-scored funnel with EBIT models and executive sign-off	<input type="checkbox"/>

## Score-Based Clinic Recommendations

**Score 6-10**

**Clinic #1**

Strategic foundations are missing. Start with **Vision to Advantage** – define your strategic question and competitive moats before any architecture work begins.

**Score 11-15**

**Clinic #2**

Strategy exists, but execution is fragmented. **Platform Powerplay** unifies scattered pilots into a coherent, scalable architecture with embedded governance.

**Score 16-18**

**Clinic #3**

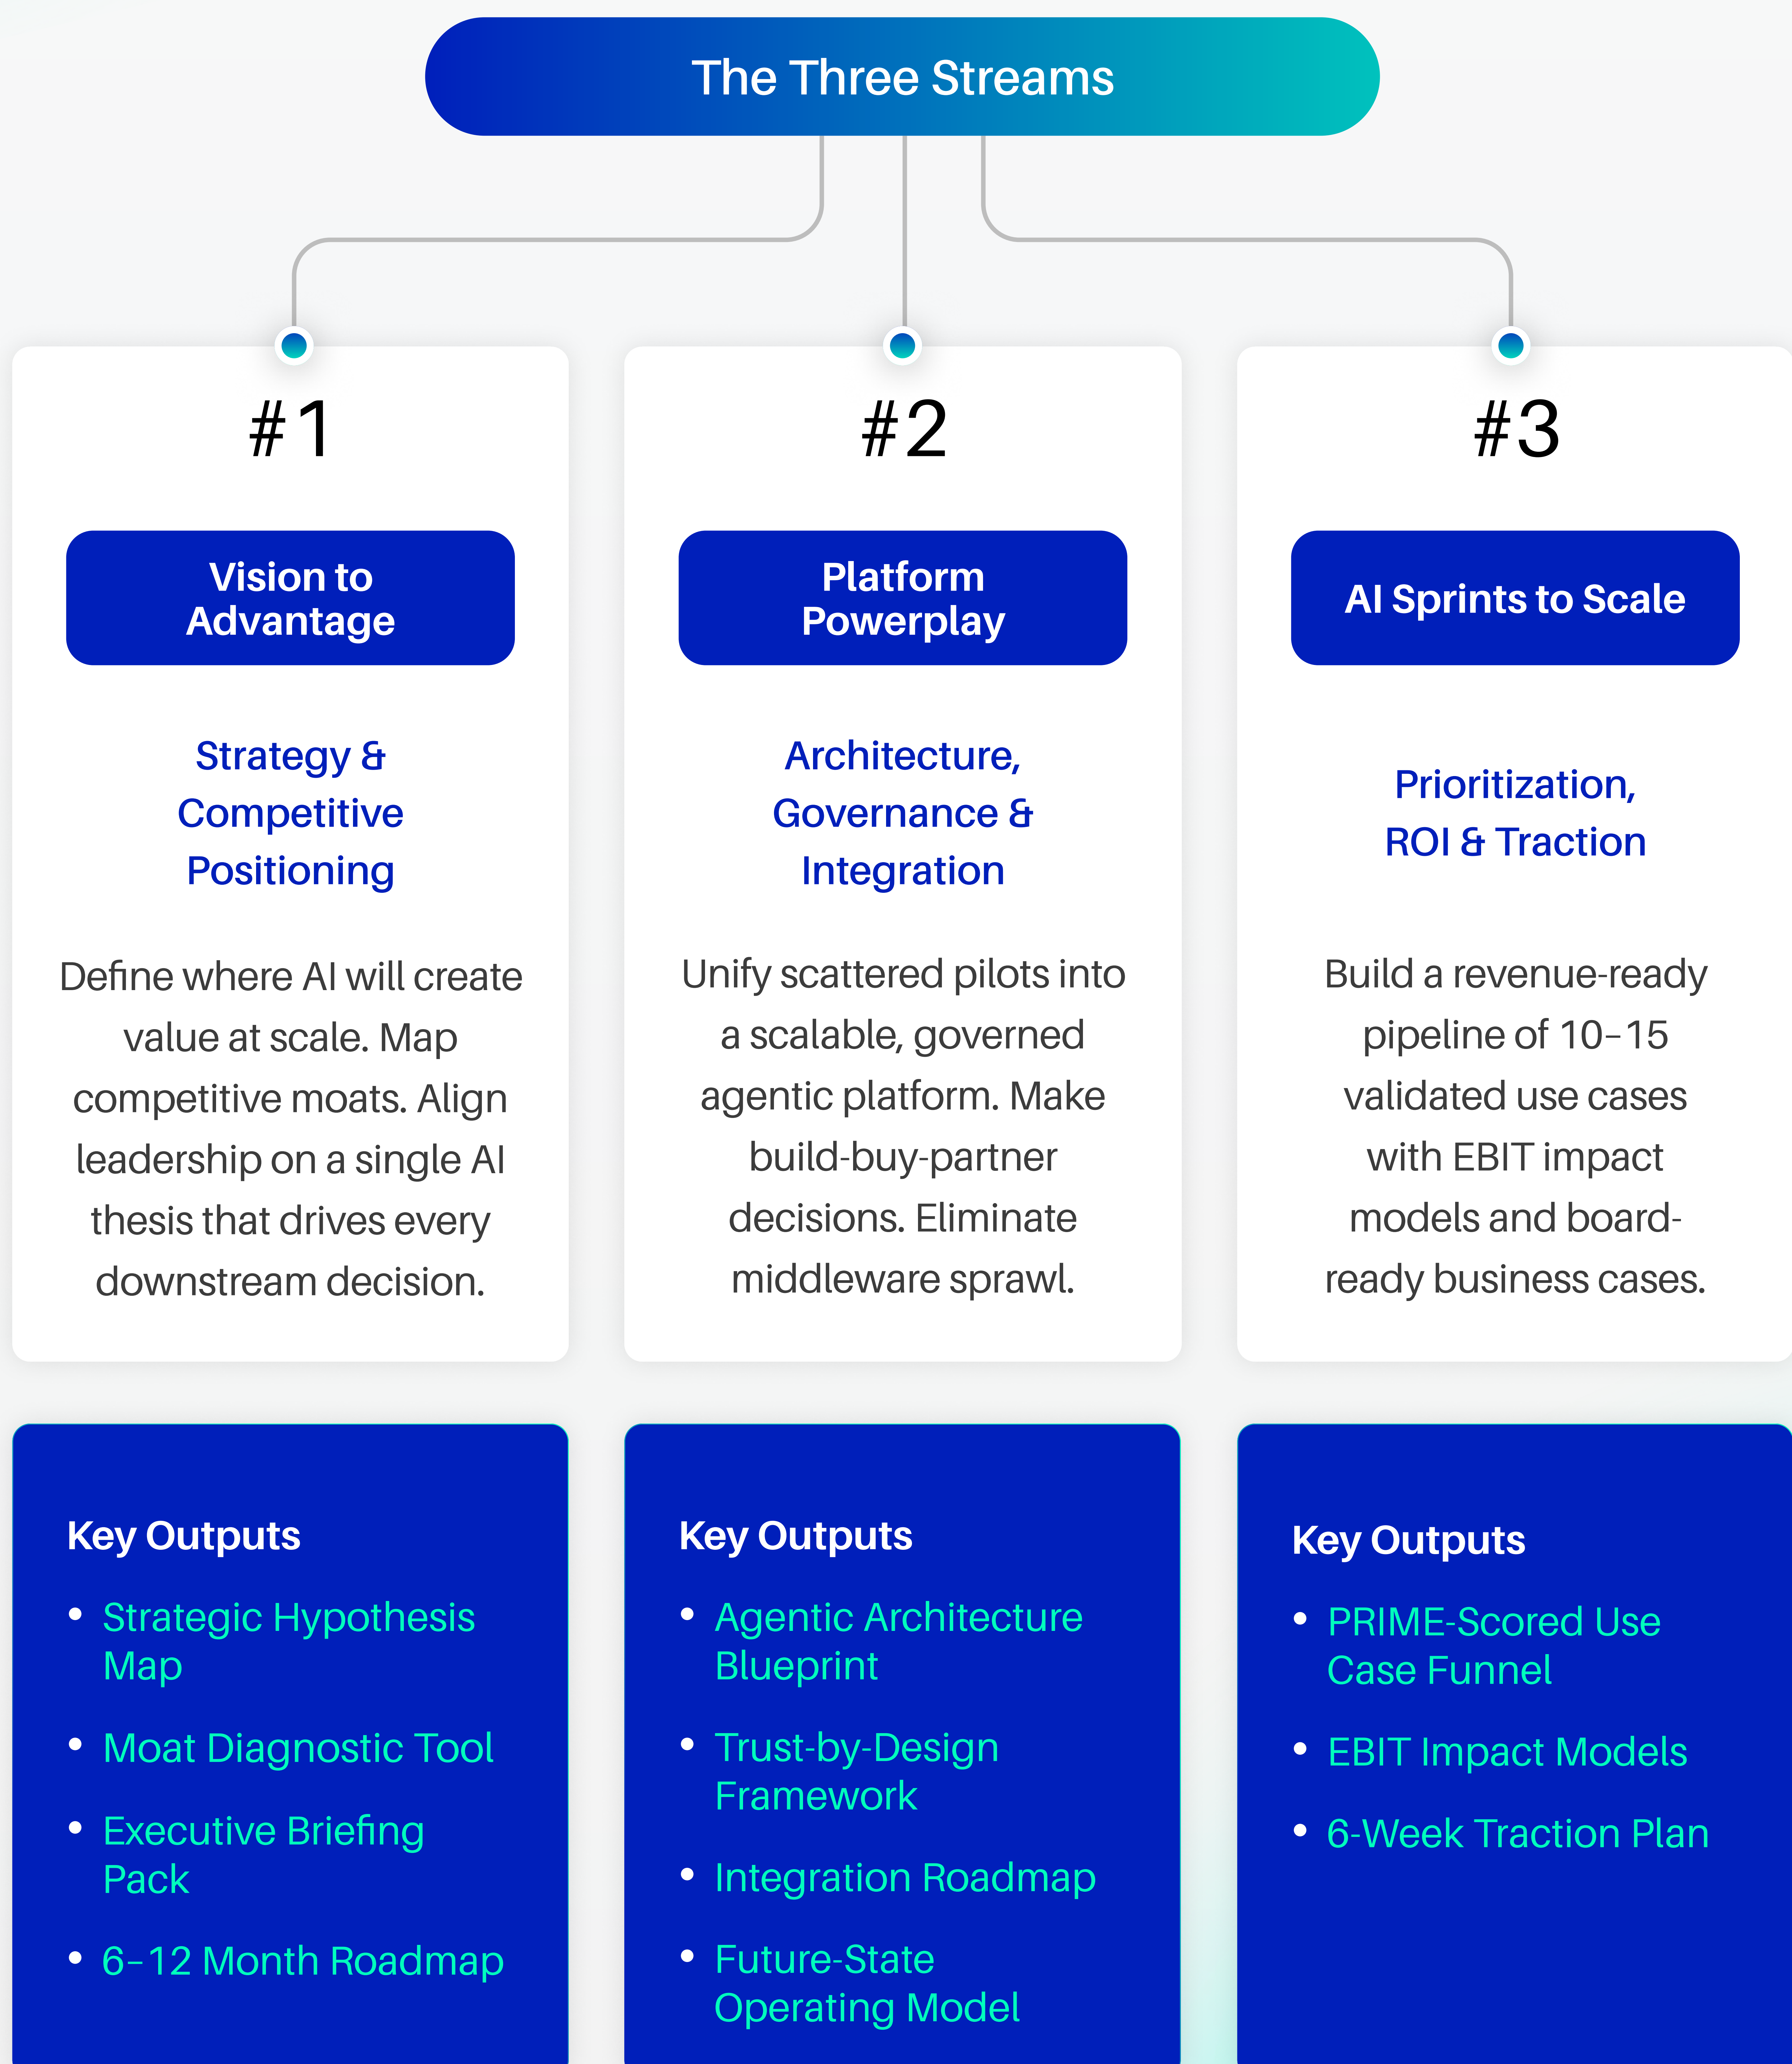
The platform is ready. Focus on **AI Sprints to Scale**: rapid prioritization of high-ROI use cases with EBIT-backed business cases and prototype validation.

**Note:** Tracks are not mutually exclusive. Many organizations benefit from combining all three Clinics in a single 3-week cycle. Pacing is calibrated to your readiness and constraints during Week 1 Discovery.



# The Solution.

## A Decision Accelerator, Not a Workshop.

The AI Strategy Clinic is a time-compressed engagement designed to produce decision clarity. Where internal steering groups debate for quarters, the Clinic produces decision-grade outputs in weeks: combining executive facilitation, technical architecture expertise, and financial modeling in a single structured engagement.



# The 3-Week Mechanism - Time-Compression Model



## WEEK 1

### Discovery & Preparation

Review strategy, current initiatives, architecture, and competitive landscape. Identify what is blocking production and agree on the core questions the Clinic will answer.

- ▶ Stakeholder interviews
- ▶ Current-state audit
- ▶ Data landscape review

## WEEK 2

### Leadership Clinic

2-3 days of structured C-suite sessions. Align on one AI thesis. Define platform architecture direction. Commit to top 3-5 bets. Kill the rest. No hedging allowed.

- ▶ Executive facilitation
- ▶ Architecture labs
- ▶ Decision gates
- ▶ Handoff to delivery

## WEEK 3

### Synthesis

Package all decisions into board deck, platform blueprint, and execution roadmap. Deliverables delivered within 1-2 weeks of the Clinic's conclusion.

- ▶ Board-ready artifacts
- ▶ EBIT models

**The Closing Commitment:** By the end of Week 3, you have a Board-ready AI narrative, a platform direction, and a prioritized pipeline with clear ownership, readiness gates, and a named executive accountable for each funded use case.



## Deep Dive - Clinic #1

# Vision to Advantage: Defining Your Strategic Question


Before any architecture or use case work is credible, leadership must answer: Where will AI create - or destroy - the most material value in our industry over the next three to five years? This is not a product roadmap question. It is a competitive positioning question, and the answer determines every downstream investment decision.

## The \$1B Question Canvas




### STEP 1 - Break-the-firm Exercise

"If an AI-native competitor entered your market tomorrow with unlimited capital and your data where would they attack first?" This surfaces strategic vulnerabilities that polite internal planning avoids.



### STEP 2 - Future-Back scenario planning

Start from a hypothetical 2028 "Agentic Future" for your sector. Work backward to identify: what capabilities must be operational by 2026 to avoid commoditization? This prevents incremental thinking disguised as ambition.



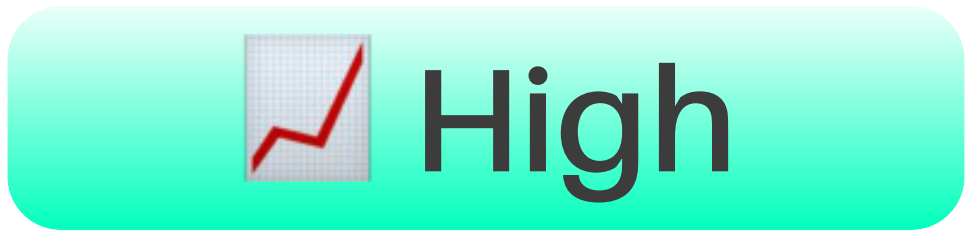




### STEP 3 - Thesis articulation

The output is a clear AI strategy thesis: where AI will create value, what outcome it will drive, and why your organization has an advantage competitors cannot easily copy.

## Moat Diagnostic - Durability in an Agentic Economy

The clinics utilize a Moat Diagnostic Tool to evaluate the strength and defensibility of a client's current competitive advantages, while also identifying where those advantages are vulnerable to erosion from emerging AI trends.

A moat is an advantage competitors cannot easily copy, buy, or build around. This diagnostic helps identify which advantages are durable, and which may weaken as AI-native competitors emerge.

Moat Type	Durability	Description	Score
Proprietary Data Flywheel	 High	Unique data that improves with use	<input type="checkbox"/> Yes <input type="checkbox"/> Partial <input type="checkbox"/> No
Orchestration Layer Control	 High	Owns the agent coordination system	<input type="checkbox"/> Yes <input type="checkbox"/> Partial <input type="checkbox"/> No
Deep Context Memory	 Medium	Institutional knowledge that competitors lack	<input type="checkbox"/> Yes <input type="checkbox"/> Partial <input type="checkbox"/> No
Distribution / Customer Access	 Medium	Embedded touchpoints AI can automate	<input type="checkbox"/> Yes <input type="checkbox"/> Partial <input type="checkbox"/> No
Regulatory / License Advantage	 Eroding	Barriers that competitors cannot purchase	<input type="checkbox"/> Yes <input type="checkbox"/> Partial <input type="checkbox"/> No

**Key insight:** Moats built on proprietary data and orchestration control are the most defensible in an agentic economy. Brand and distribution are necessary but increasingly insufficient on their own.

## Primary Deliverable: The Executive Briefing Pack

- **Strategic Hypothesis Match**

A structured set of testable bets defining where AI will create or destroy the most value across your business lines, giving leadership a single, aligned thesis to drive every downstream investment decision.

- **Risk Landscape Brief**

A high-level view of your regulatory exposure (EU AI Act, GDPR, sector-specific rules), competitive threats from AI-native entrants, and ethical risks that require early mitigation before build begins.

- **Transformation Roadmap V1**

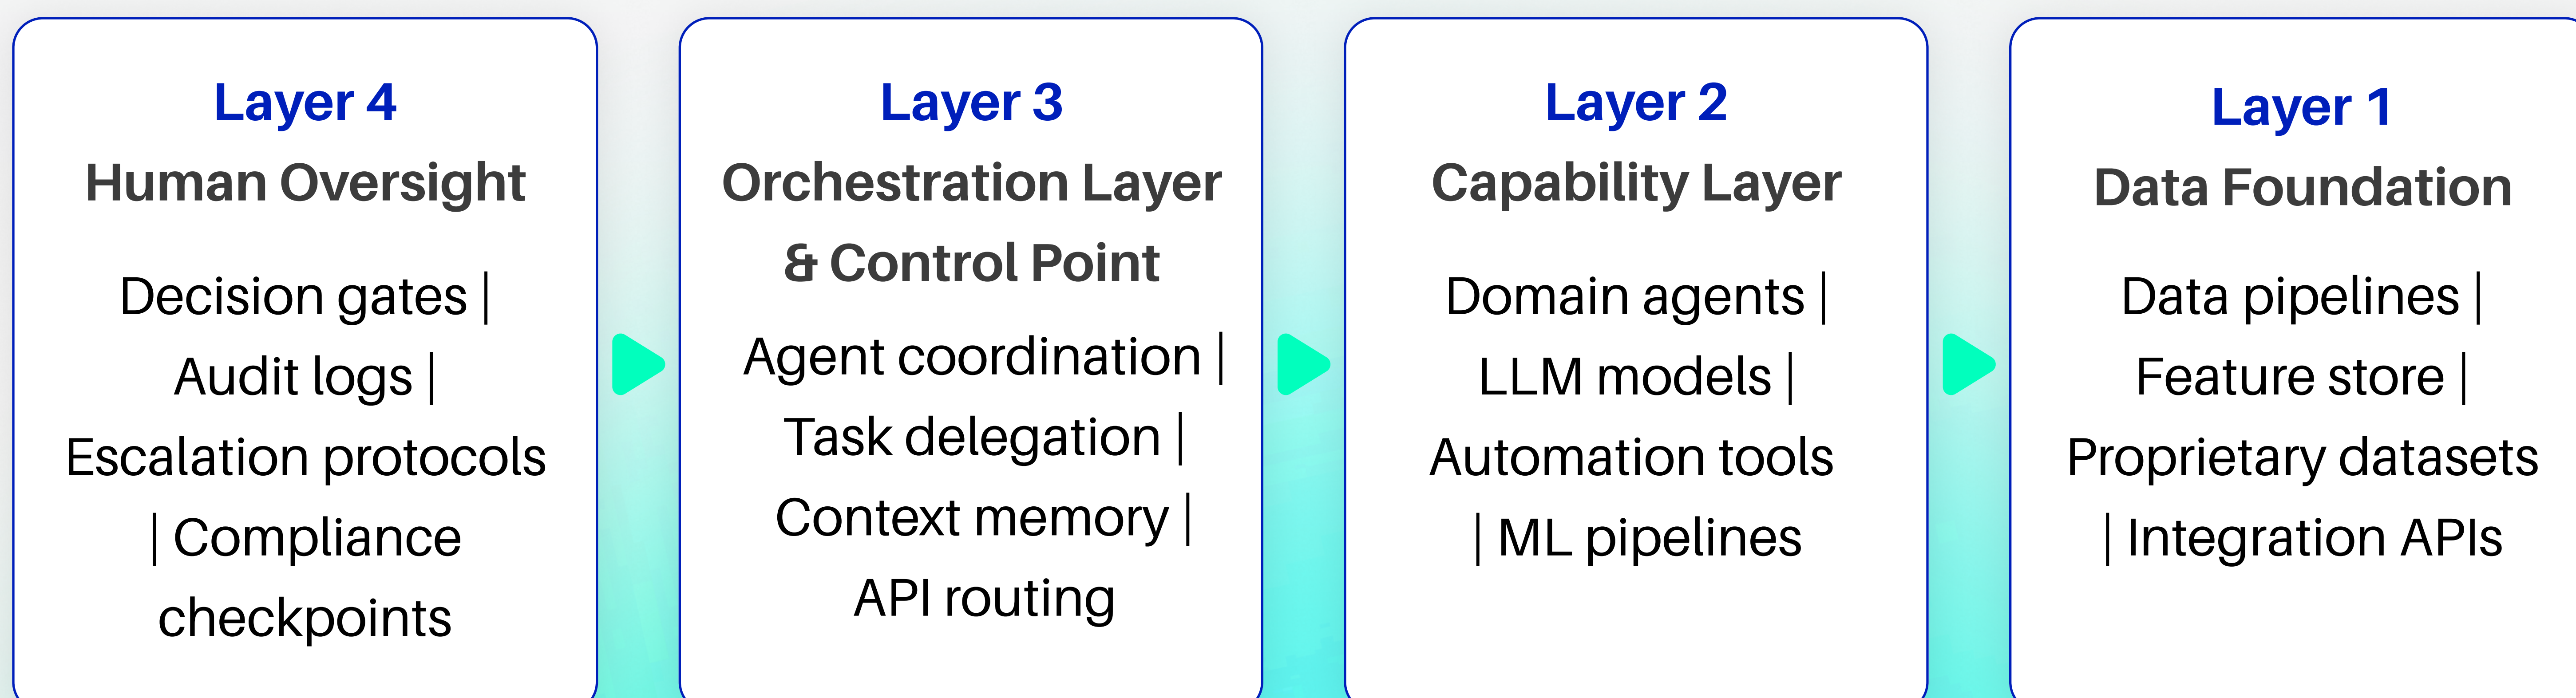
A high-level 6-12 month timeline covering capability building, governance structure, and phased execution priorities - structured for board presentation and ready to move into delivery on day one.

## Deep Dive - Clinic #2

### Platform Powerplay: From Scattered Pilots to Unified Architecture

Most enterprise AI programs struggle to scale because they treat architecture as an engineering concern rather than a strategic one. The Orchestration Layer - the system that coordinates how agents interact, share context, and escalate to humans - is not a commodity. It is a competitive control point that determines how fast new use cases can be deployed and how governable the system remains at scale.

### Agentic Architecture Blueprint - Layer Model



## Build-Buy-Partner Strategy Canvas

The Build-Buy-Partner Strategy Canvas is a structured decision-making framework utilized in Clinic 2 (Platform Powerplay). Its primary purpose is to help executive and technical leadership make explicit, strategic sourcing decisions regarding AI platform components and capabilities.

By systematically evaluating whether to develop a solution internally, purchase it off-the-shelf, or collaborate with a partner, the canvas helps organizations align their investments, balance speed-to-market with strategic control, and reduce the risks of vendor lock-in and fragmented integration

Component	Build	Buy	Partner
Orchestration Layer	 Recommended	 Risk: Lock-in	 Partial
Foundational LLMs	 Rarely	 Recommended	 Optional
Domain-Specific Agents	 Recommended	 Evaluate	 Recommended
Data Pipelines	 Partial	 Recommended	 Recommended
Governance & Audit Tools	 Partial	 Recommended	 Optional

Ultimately, this feeds into a set of Strategic Recommendations and a Risk & Dependency Map, ensuring procurement and transformation teams have a clear mandate on what to acquire versus what to build.

## Trust-by-design Framework

Governance is an architectural decision made at design time, not a compliance checkbox added at launch. Four embedded controls: human-in-the-loop gates on high-consequence actions, audit-log-by-default architecture, role-based decision rights for agent actions, and EU AI Act risk tier classification for each use case at intake.

## Integration Roadmapping - Illustration Sequence

### Foundation (Months 1-3)



Data pipelines, API governance, Governance model

### Orchestration (Months 3-7)



Orchestration layer, First agents, Internal pilot

### Expansion (Months 7-12)



Multi-agent systems, Change management, External deployment

# Deep Dive - Clinic #3

## AI Sprints to Scale: A Revenue-Ready Use Case Pipeline

The challenge is not generating AI ideas, but to kill the wrong ones before they consume resources. The Sprints to Scale track introduces a structured triage mechanism to rank, validate, and financially model 10-15 use case candidates, producing a Board-ready portfolio with clear fund/kill decisions.

### The PRIME Framework - Use Case Evaluation Scoring

Each use case candidate is scored 1-5 across five dimensions. The weighted total determines its position in the funded pipeline. Low PRIME scores trigger a kill decision - not a reschedule.

#### **P** Potential (Score / 5):

Maximum addressable EBIT impact at full production scale - quantified as revenue uplift, cost reduction, or risk avoidance. (1 = < \$1M 5 = > \$50M annual impact.)

#### **R** Reach (Score / 5):

Number of users, transactions, or processes affected. Enterprise-wide automation scores higher than single-function tools. (1 = one team 5 = cross-enterprise or customer-facing.)

#### **I** Impact on Moat (Score / 5):

Does this use case build, protect, or just maintain competitive advantage? Moat-building cases score highest. (1 = efficiency only. 5 = creates new durable moat.)

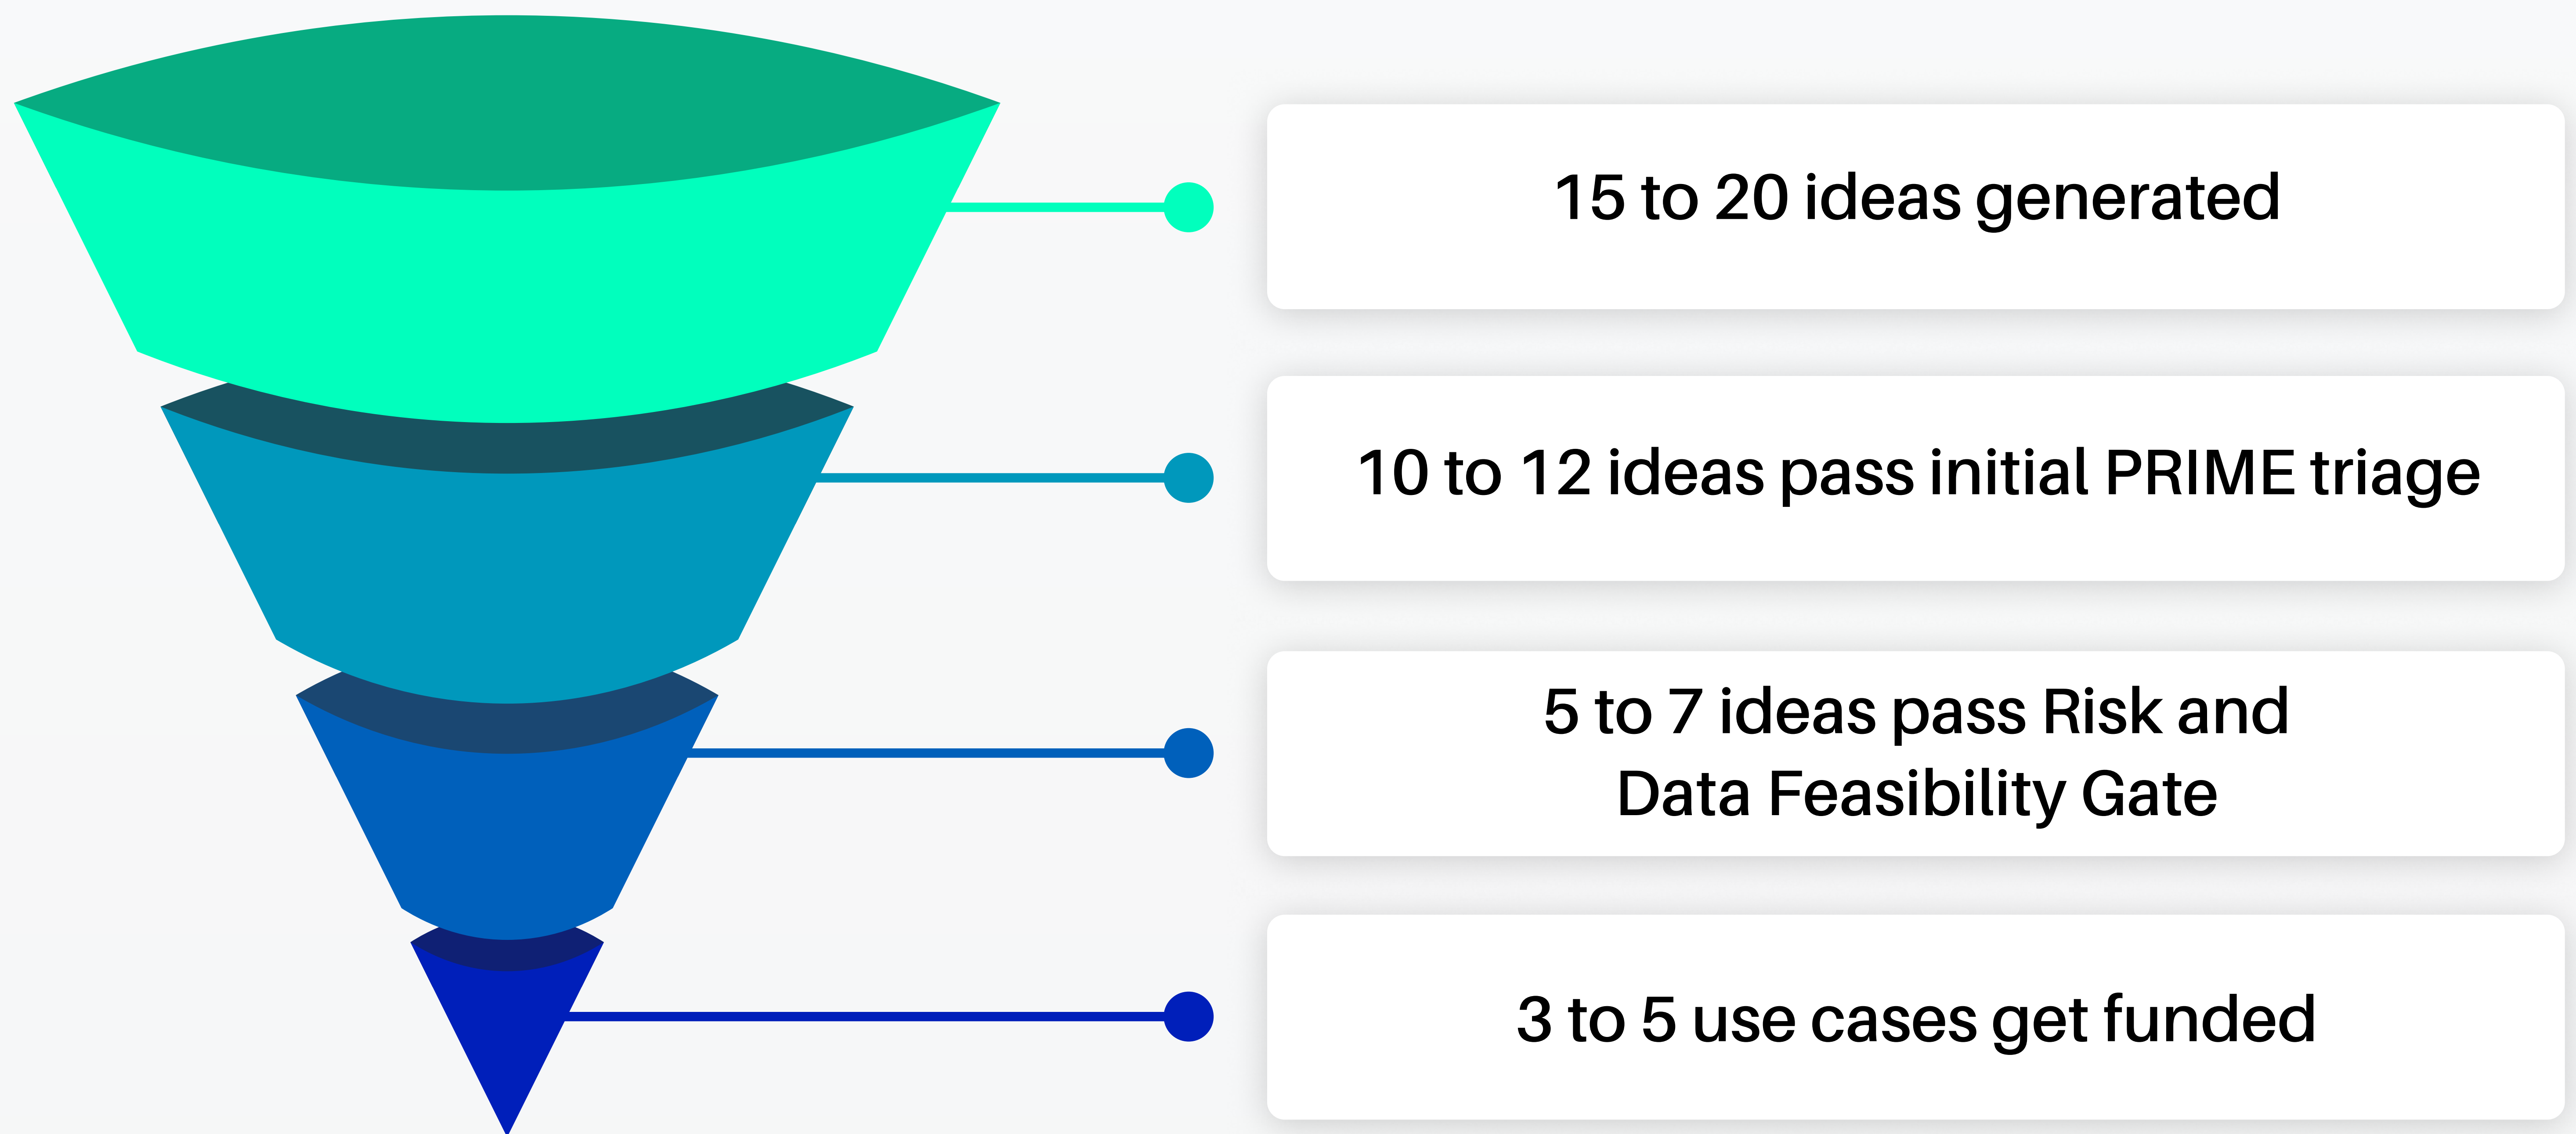
#### **M** Measurability (Score / 5):

Can success be measured unambiguously within 90 days of deployment? Clear leading and lagging KPIs are required for funding approval. (1 = vague outcomes · 5 = automated, real-time P&L attribution.)

#### **E** Effort & Feasibility (Score / 5, inverted):

Total cost of delivery including data readiness, engineering complexity, and change management. Lower effort scores higher. Gating risks (compliance, data availability) evaluated here. (1 = unresolved blockers · 5 = ready to build now.)

## Prioritization Funnel



## Estimated EBIT Impact Model - Per Use Case

Use Case	PRIME Score	Year 1 EBIT	Payback period	Decision
Use Case A	● 22/25	▲ High (over \$5m)	🚀 Under 12 months	✅ Proceed
Use Case B	● 19/25	■ Medium (\$1m to \$5m)	🕒 Around 18 months	⚠️ Proceed, confirm assumptions
Use Case C	● 14/25	▼ Low (under \$1m)	🐢 Around 24 months	⚠️ Only if strategic
<b>Below Threshold</b>	● <10/25	■ Not worth pursuing	☐ Not applicable	⊘ Stop

**Note:** Example model for reference only. Any projected impact should be reviewed and validated internally before making investment decisions.

# Commitment & Action

## Your 30/60/90 Day Execution Plan

By the end of the 3-week Clinic, this becomes your operating commitment - with named owners, funded decisions, and a first milestone that proves momentum within 90 days.

### Day 0–30: Align & Diagnose

Readiness assessment complete → Clinic track confirmed → Executive sponsor named → Week 1 Discovery session scheduled → Data Feasibility Gate criteria agreed.

### Day 31–60: Decide & Commit

Clinic delivered. AI thesis articulated → Top 3–5 use cases PRIME-scored and ranked → Platform architecture direction set → Build/Buy/Partner decisions made → Board briefing pack drafted

### Day 61–90: Build Momentum

First funded use case in active build → 6-week traction milestone defined and tracked → Governance model operational → Executive sponsor reviewing weekly KPI report


## Next Steps

### A Strategic Alignment Briefing


Use this session to pressure-test your readiness score, confirm which Clinic track fits your current priorities, and outline a tailored engagement structure.

You leave with a clearer view of where to start and how to move into the first Clinic session with a defined mandate.

## Book a 30-Minute Strategy Briefing



Fixed-fee, no variable cost surprises



Board-ready deliverables in 15 days



Optional credit-back when you proceed to build with Cikulum.

