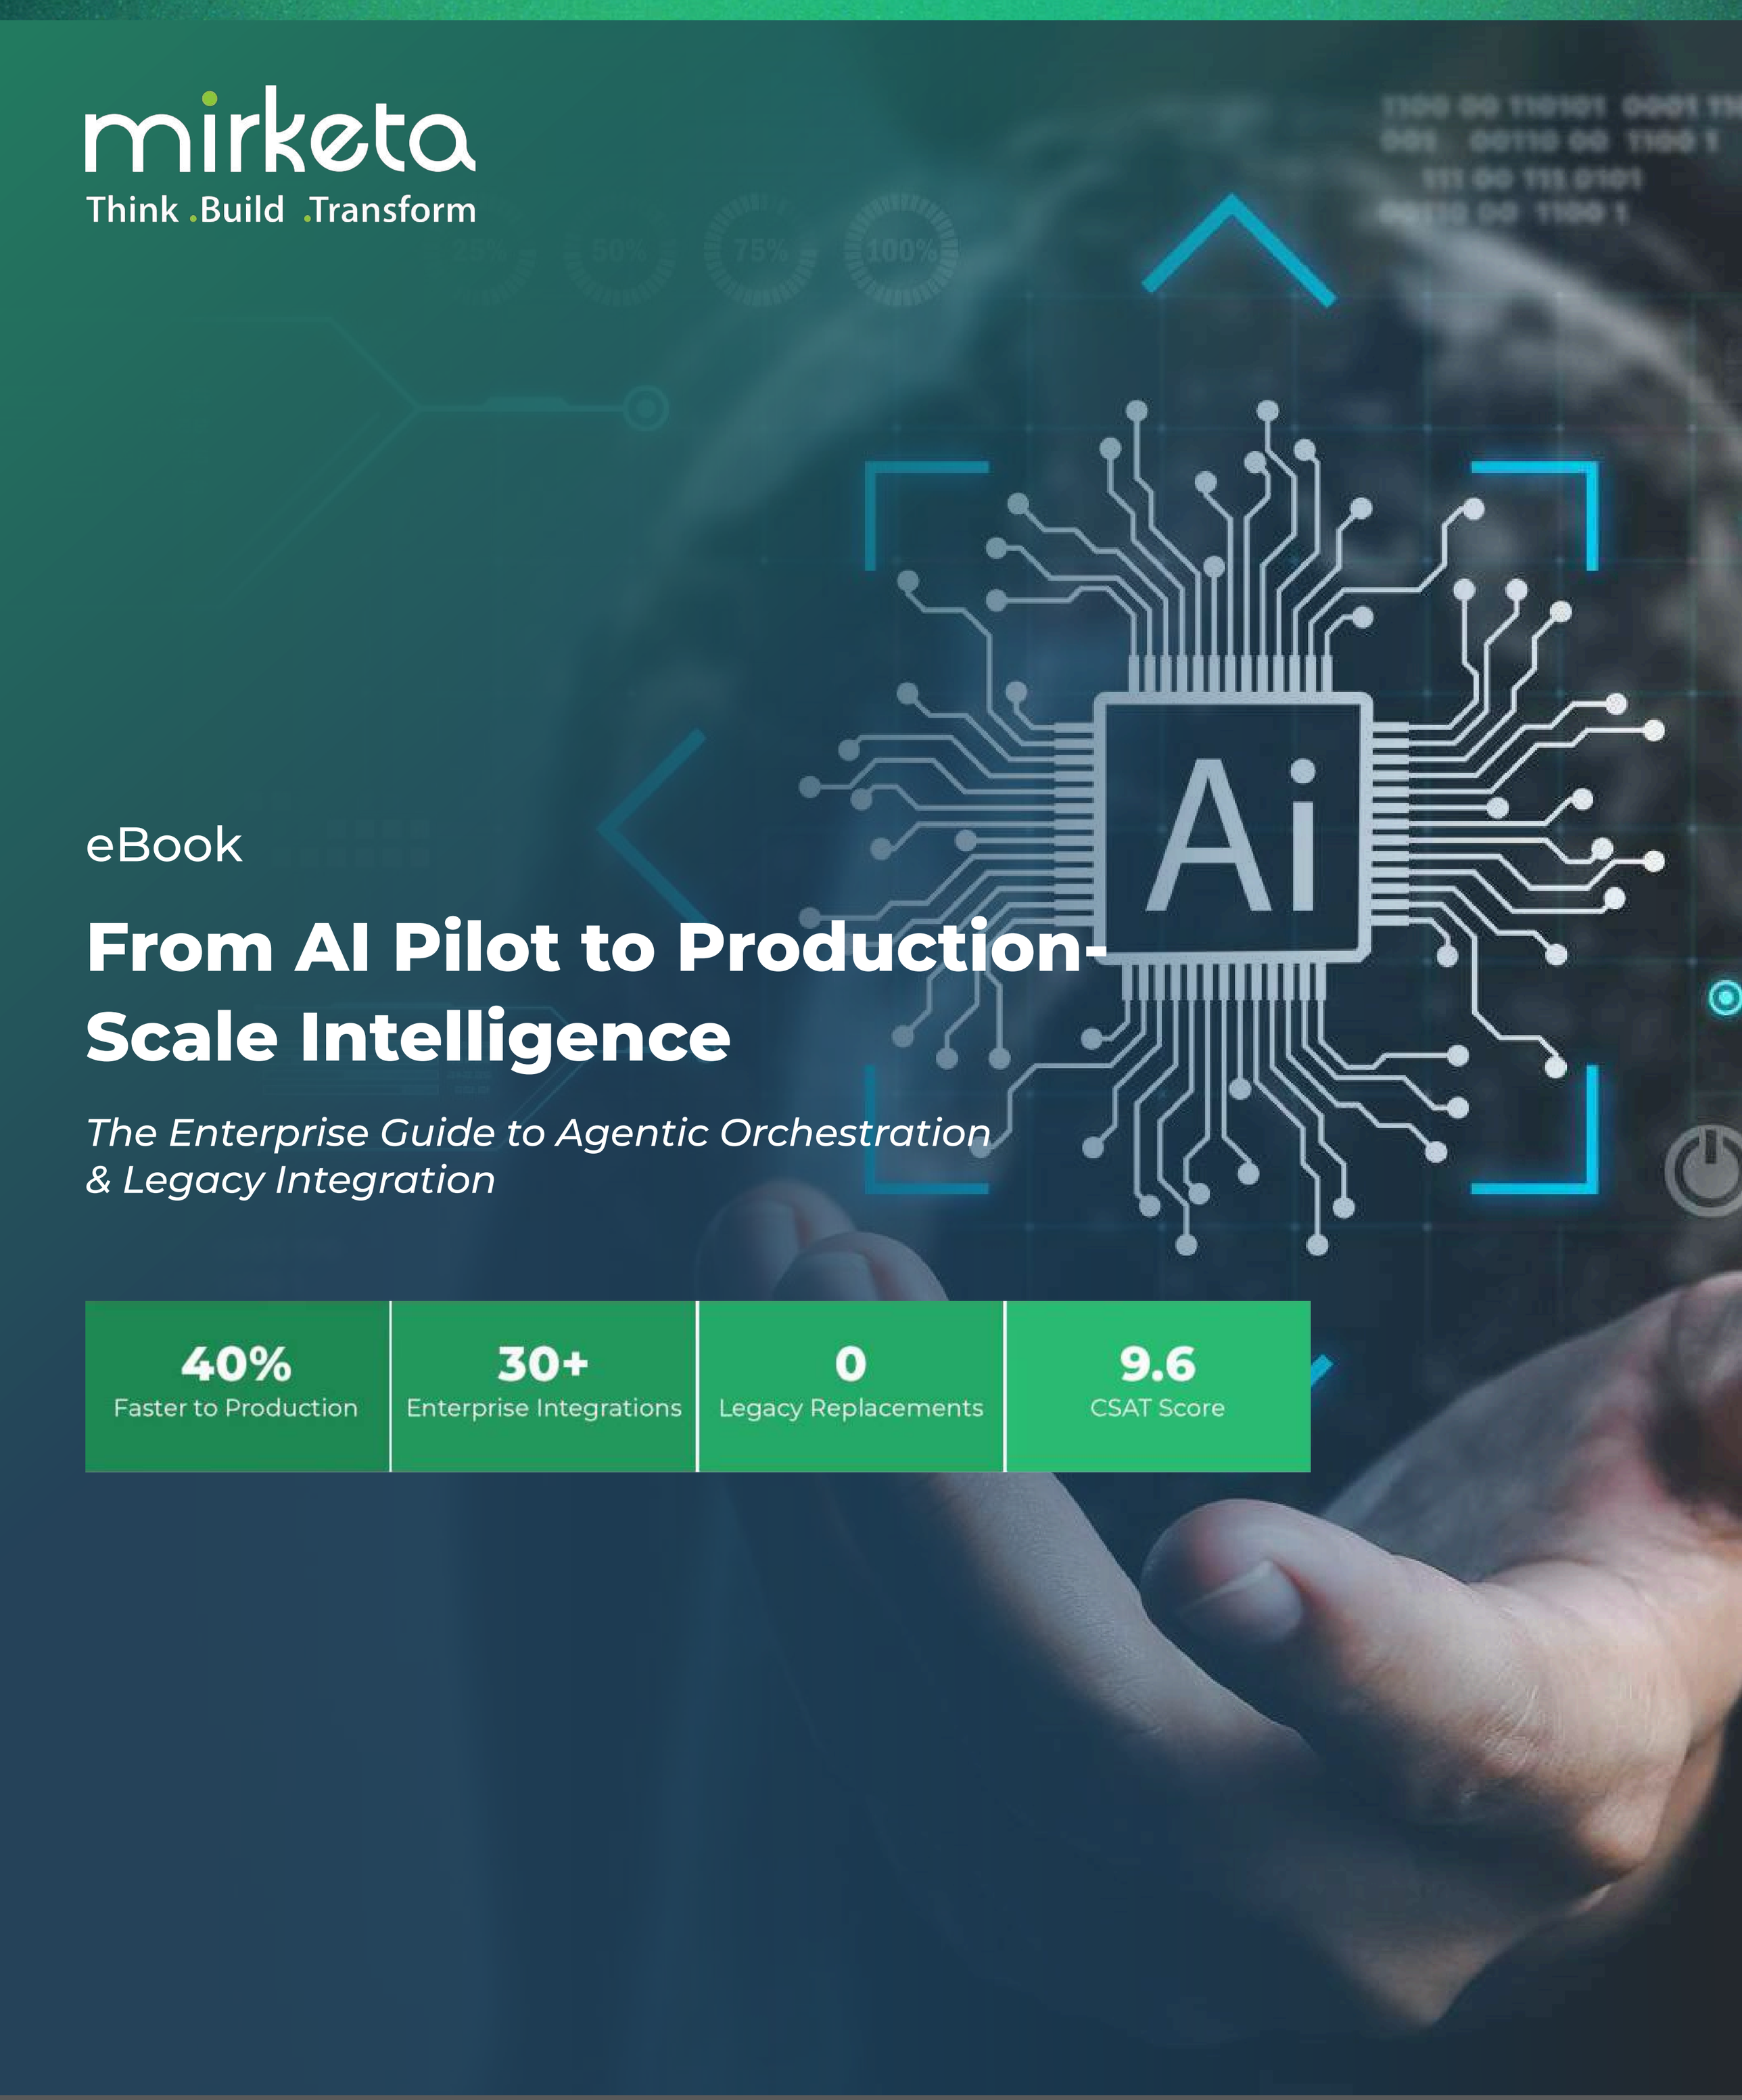
mirketa

Think .Build .Transform

eBook

From AI Pilot to Production-Scale Intelligence

The Enterprise Guide to Agentic Orchestration & Legacy Integration



40%	30+	0	9.6
Faster to Production	Enterprise Integrations	Legacy Replacements	CSAT Score

Table of Contents

01 **Introduction: The AI Stall Problem** 02

02 **Chapter 1: Why AI Initiatives Stall Before They Scale** 02

The M X N Integration Trap

The Legacy System Misconception

03 **Chapter 2: The Orchestration Solution** 03

What Is AI Orchestration?

Model Context Protocol (MCP): The Universal Integration Standard

Multi-Agent Orchestration: Beyond Single-Model AI

04 **Chapter 3: Three Architectural Patterns For Production AI** 04

Pattern 1: Orchestrator-Subagent

Pattern 2: Parallel Fan-Out

Pattern 3: Directed Acyclic Graphs (DAG)

05 **Chapter 4: The Trust & Security Layer** 05

06 **Chapter 5: Industry Applications** 06

Healthcare

Nonprofit & Education

Technology & SaaS

07 **Chapter 6: The Delivery Framework** 07

08 **Chapter 7: Mirketa Vs. The Alternatives** 08

09 **Conclusion: Your Pilot Deserves To Become Production** 09

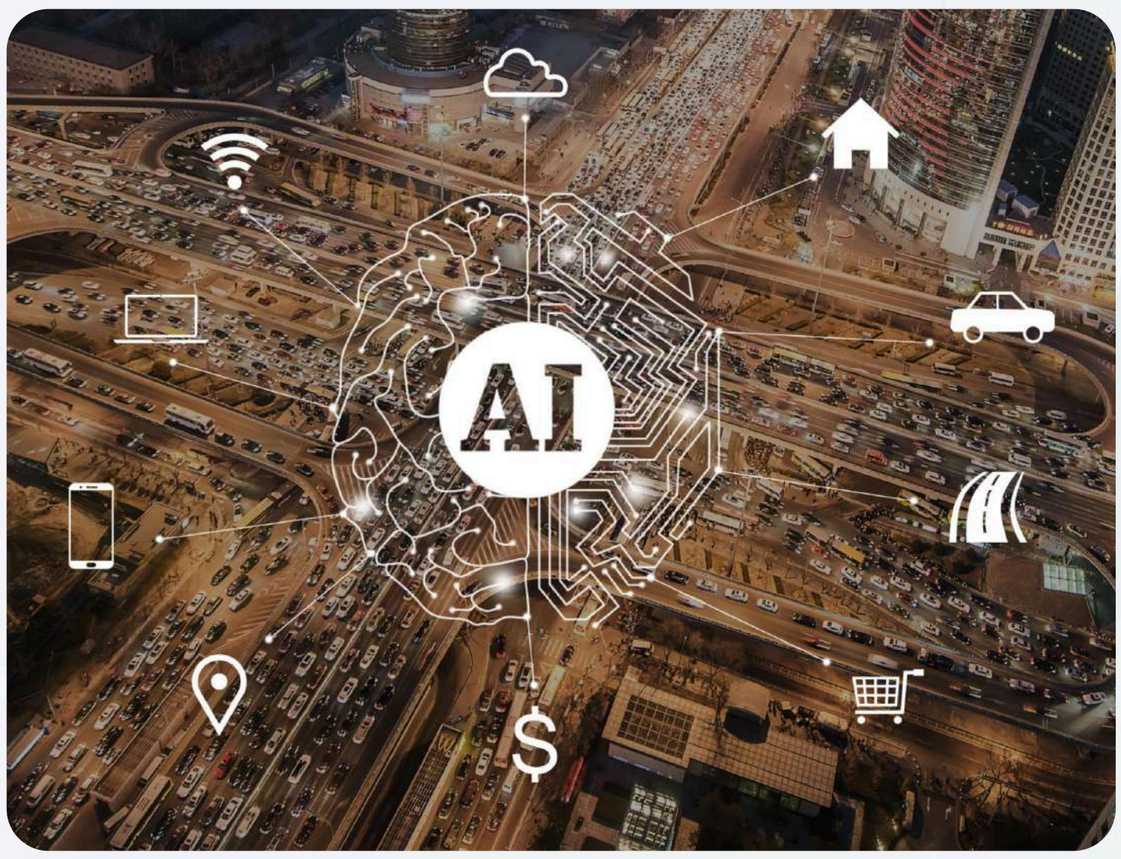
Introduction: The AI Stall Problem

Most enterprises are not short on ambition when it comes to artificial intelligence. In boardrooms across healthcare, financial services, manufacturing, and technology, AI pilots are greenlit, budgets are approved, and proof-of-concept projects are launched with fanfare.

Then they stall.

Not because the technology fails. Not because the use case was wrong. But because the path from a contained pilot to a production-grade, enterprise-wide system turns out to be far harder than anticipated. The real bottleneck is not intelligence, it is orchestration.

This ebook explores why most AI initiatives stall before they scale, what the right architectural approach looks like, and how Mirketa's AI Orchestration framework helps enterprises move from isolated AI experiments to unified, production-scale intelligence, without ever replacing the legacy systems they already depend on.



Chapter 1: Why AI Initiatives Stall Before They Scale

The M x N Integration Trap

Here is a pattern that plays out in enterprise after enterprise: A team launches an AI initiative. They connect a language model to one data source via a bespoke adapter. Success. Then they want to connect a second model to two data sources. Then a third model, four data sources, and a third-party tool.

Suddenly, instead of a clean system, they have an M x N integration mess, every new model or data source requires a hand-crafted connector to every other. The complexity grows exponentially. Context windows overflow, degrading AI response quality. Governance becomes impossible. And the production deployment that was supposed to follow the pilot keeps getting pushed back.



The M x N Integration Trap

- Bespoke adapters multiply with every new model or data source
- AI context windows overflow, causing degraded response quality
- No consistent governance across agents and tools
- Legacy systems are locked out of AI value creation
- Pilot projects cannot be promoted to production environments

The Legacy System Misconception

A common, costly misconception: that legacy systems must be replaced before an organization can participate in the AI era. This belief drives 18-month migration projects, ballooning budgets, and enormous disruption, often before a single AI agent goes live.

The truth is the opposite. Your legacy systems contain decades of business-critical data, institutional knowledge encoded in workflows, and relationships with your most important customers. The goal is not to discard this value, it is to unlock it for AI.



Mirketa has integrated 30+ enterprise systems across every major industry without ever requiring a client to replace their existing infrastructure. The systems you already own are not obstacles to AI. With the right orchestration layer, they become its foundation.

Chapter 2: The Orchestration Solution

What is AI Orchestration?

AI orchestration is the coordinated management of multiple AI agents, data sources, and enterprise systems to achieve complex business goals reliably and at scale. Without orchestration, AI initiatives remain isolated pilots, each model operating independently, unable to access the full context of business data or trigger actions across systems.

Orchestration is what transforms individual AI capabilities into an integrated, production-grade intelligence layer that actually runs your business processes. It is the difference between a smart demo and a system your operations can depend on.

Model Context Protocol (MCP): The Universal Integration Standard

At the heart of Mirketa's orchestration approach is the Model Context Protocol (MCP), an open standard developed by Anthropic that provides a universal interface between AI models and external data sources and tools.

Instead of building custom integrations for every system AI needs to access, MCP creates a single standardized interface. Mirketa uses MCP to expose enterprise systems, CRM, ERP, legacy databases, and APIs as AI-accessible Resources (for reading context) and Tools (for executing actions). This eliminates the M x N complexity permanently and makes an organization's entire data estate available to AI agents through one governed layer.

MCP Technical Capabilities

- Transport-agnostic: works via stdio, HTTP, and SSE protocols
- Exposes enterprise data as secure, AI-readable Resources
- Enables executable Tools for cross-system workflow automation
- Single integration layer replaces dozens of custom adapters
- Eliminates context bloat with managed context windows

Multi-Agent Orchestration: Beyond Single-Model AI

True enterprise AI is not a single model answering questions. It is a coordinated system of specialist agents, each optimized for a specific function, working together to accomplish complex, multi-step workflows. Mirketa's orchestration framework coordinates multiple AI agents, memory stores, and human touchpoints into reliable, long-running workflows.

Complex tasks are decomposed and distributed across specialist agents for maximum precision. The result: AI that does not just respond to queries but actively drives business processes from end to end.

3.4 Operations and Finance

Expense Categorization and Anomaly Detection
AI can automatically categorize transactions, flag anomalies, and assist with budget forecasting — reducing the manual burden on your finance team and improving accuracy.

Knowledge Management
AI-powered knowledge bases can surface relevant policies, procedures, and past decisions instantly — reducing onboarding time and preserving institutional knowledge even when staff turn over.

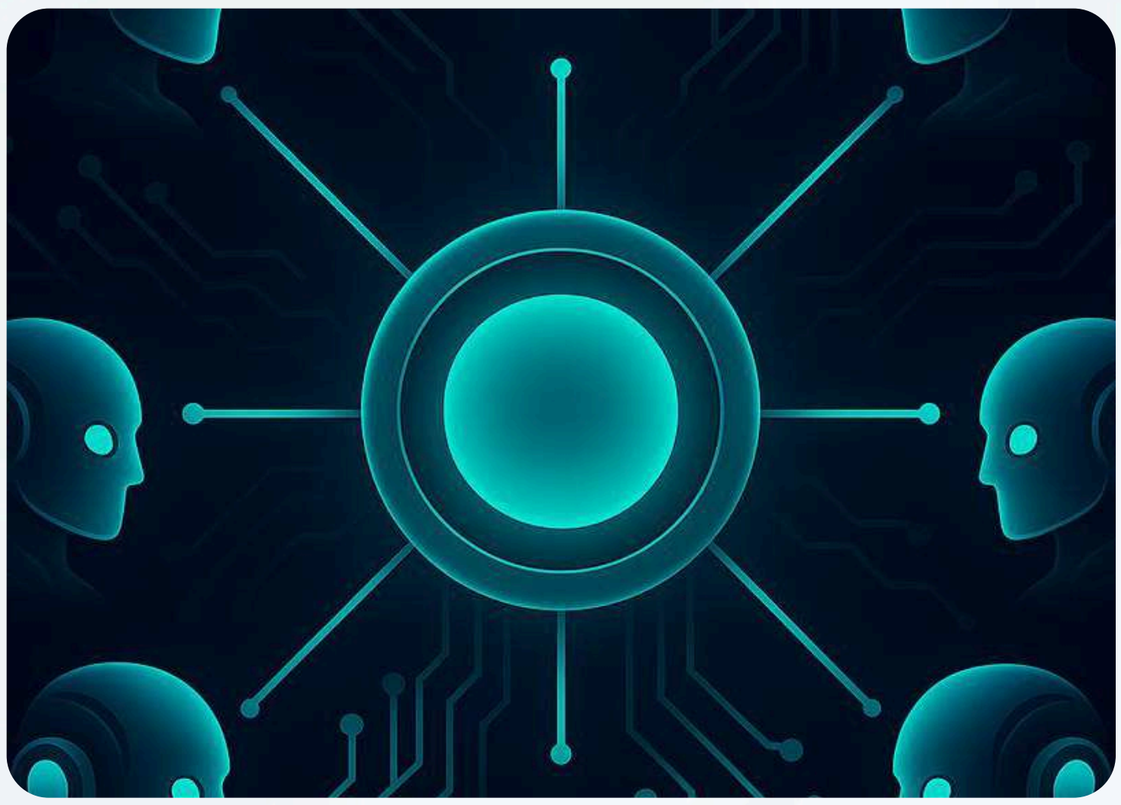
Chapter 3: Three Architectural Patterns for Production AI

Mirketa deploys three proven orchestration patterns, each selected based on workflow complexity and performance requirements. The right architecture is not one-size-fits-all, it depends on what your specific business processes demand.

Pattern 1: Orchestrator-Subagent

A central orchestrator agent decomposes complex tasks and delegates to specialist subagents, each optimized for a specific function. Results are aggregated and validated before returning to the user, ensuring precision at every step.

Best for: Complex, multi-step workflows that require coordination across multiple systems or data sources.



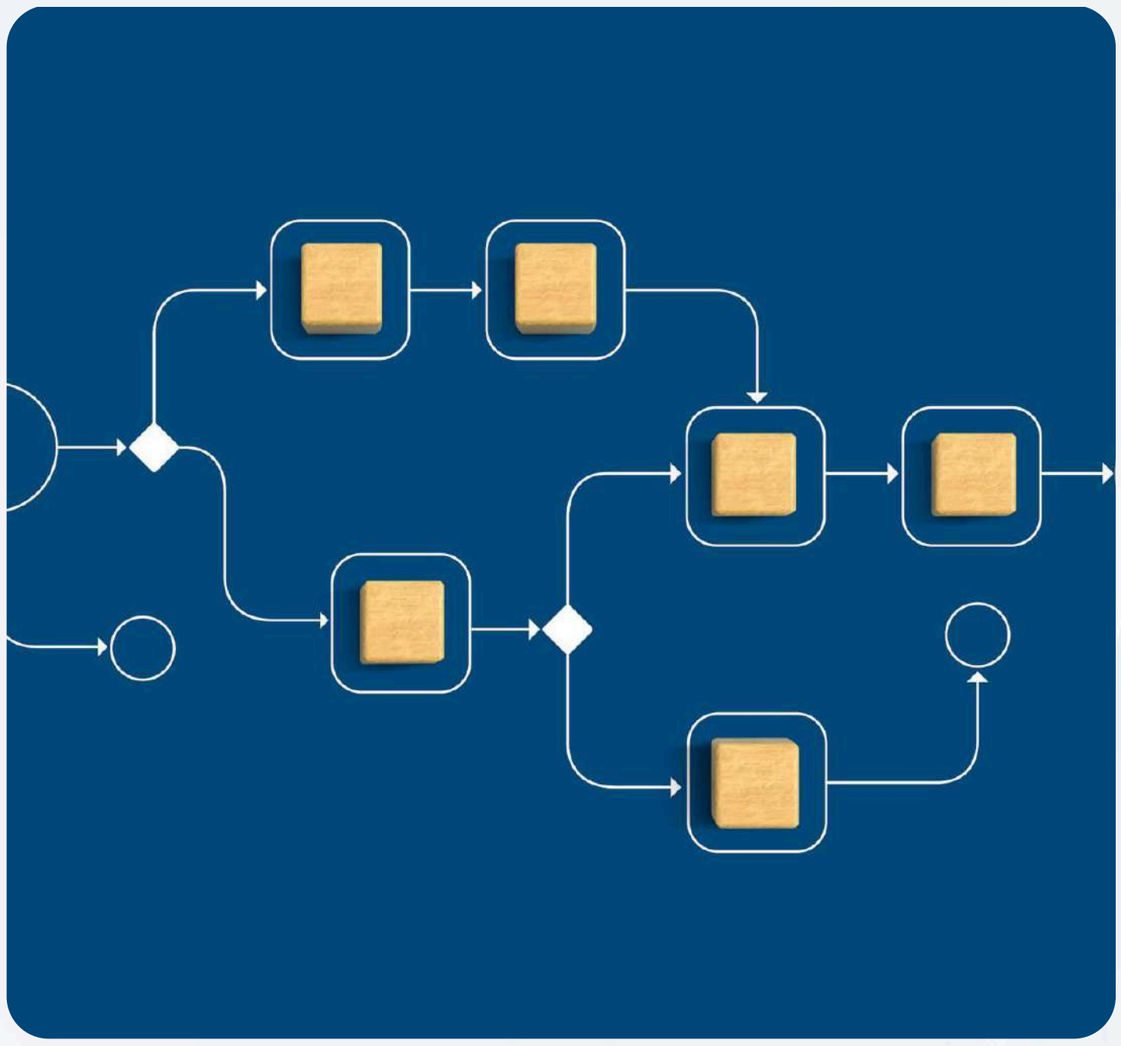
Example: A healthcare network using an orchestrator to simultaneously query the EHR system, check scheduling availability, verify insurance eligibility, and draft a care coordination note, all in a single patient interaction.

Pattern 2: Parallel Fan-Out

Multiple agents spawn simultaneously across independent data streams, dramatically reducing total latency. Ideal for scenarios where tasks can be executed in parallel, such as multi-source data retrieval or concurrent analysis across business units.

Best for: High-volume, time-sensitive tasks where speed is critical.

Example: A financial services firm running simultaneous fraud detection analysis across multiple transaction streams, cutting triage time by 3x compared to sequential processing.





Pattern 3: Directed Acyclic Graphs (DAG)

Sequential dependencies are modeled as a DAG, ensuring each processing step enriches the next in a controlled, auditable pipeline. Every node's output becomes the next node's context, building intelligence progressively through the workflow.

Best for: Sequential, dependency-driven pipelines where auditability and compliance are paramount.

Example: A compliance automation pipeline where a policy database is queried, an AI review flags potential issues, a human approves the finding, and a final report is generated, with every step logged and auditable.

Chapter 4: The Trust & Security Layer

Enterprise AI must be trustworthy before it can be transformative. This is not a platitude, it is the practical reason most enterprise AI initiatives fail to reach production. Leadership, compliance teams, and boards will not approve autonomous AI systems that they cannot audit, control, or override.

Mirketa's Trust Layer is built into every orchestration deployment. It implements progressive oversight, starting with maximum human control and systematically reducing approval friction as the system demonstrates consistent, reliable performance.

Three Pillars of the Mirketa Trust Layer

- Approval Gates: Explicit human consent required before financial transactions, record modifications, or customer communications
- Ambiguity Resolution: When agent confidence falls below a defined threshold, it escalates to your team with structured options, never guesses
- Trust Calibration: Starts with maximum oversight and progressively reduces approval friction as performance is validated

The system also includes bounded retry logic to prevent runaway agent loops, and every action is logged for compliance and audit purposes. The result is an AI system that earns the right to operate autonomously, rather than one that demands trust before it has demonstrated reliability.

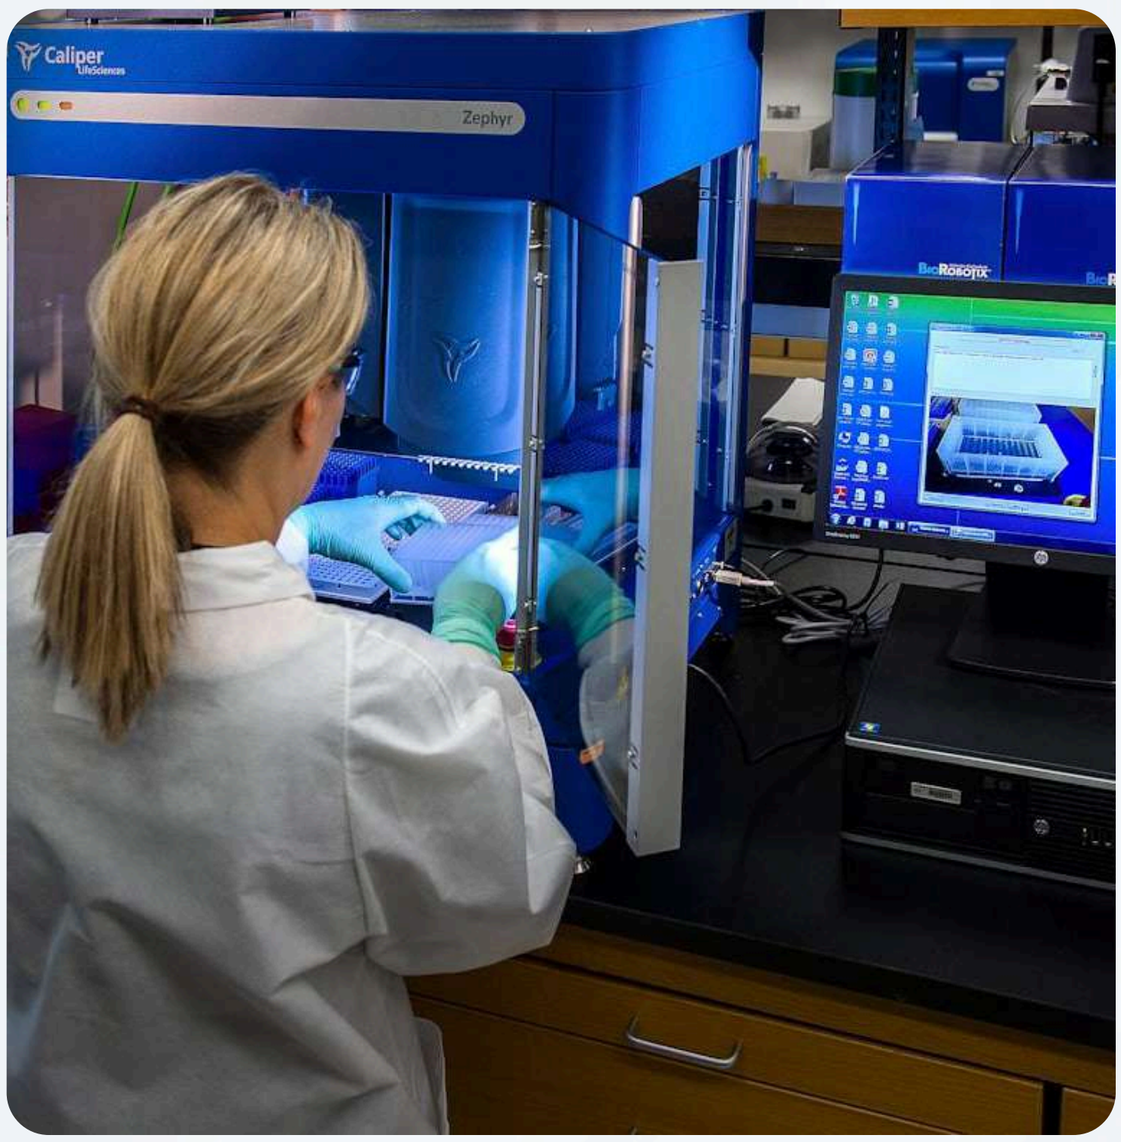
> "We were nervous about AI agents making autonomous decisions in our financial workflows. Mirketa's Approval Gates and Trust Calibration framework gave our compliance team the confidence to approve production deployment. The system started with maximum oversight and earned its autonomy over time."
>
> **— Chief Technology Officer, Regional Financial Services Firm**

Chapter 5: Industry Applications

Healthcare

Healthcare organizations manage a fragmented mix of EHR systems, payer platforms, scheduling tools, and clinical databases. Mirketa's orchestration layer bridges these silos, enabling AI agents to access patient data, trigger care workflows, and coordinate across clinical teams in real time, within HIPAA-compliant governance frameworks.

- 50%+ Faster chart review with AI Patient Chart Summaries
- 40% Reduction in call center volume via Smart Appointment Agent
- 30-40% Fewer no-shows with predictive intervention workflows
- 25% Higher first-pass claim acceptance via orchestrated pre-submission review





Nonprofit & Education

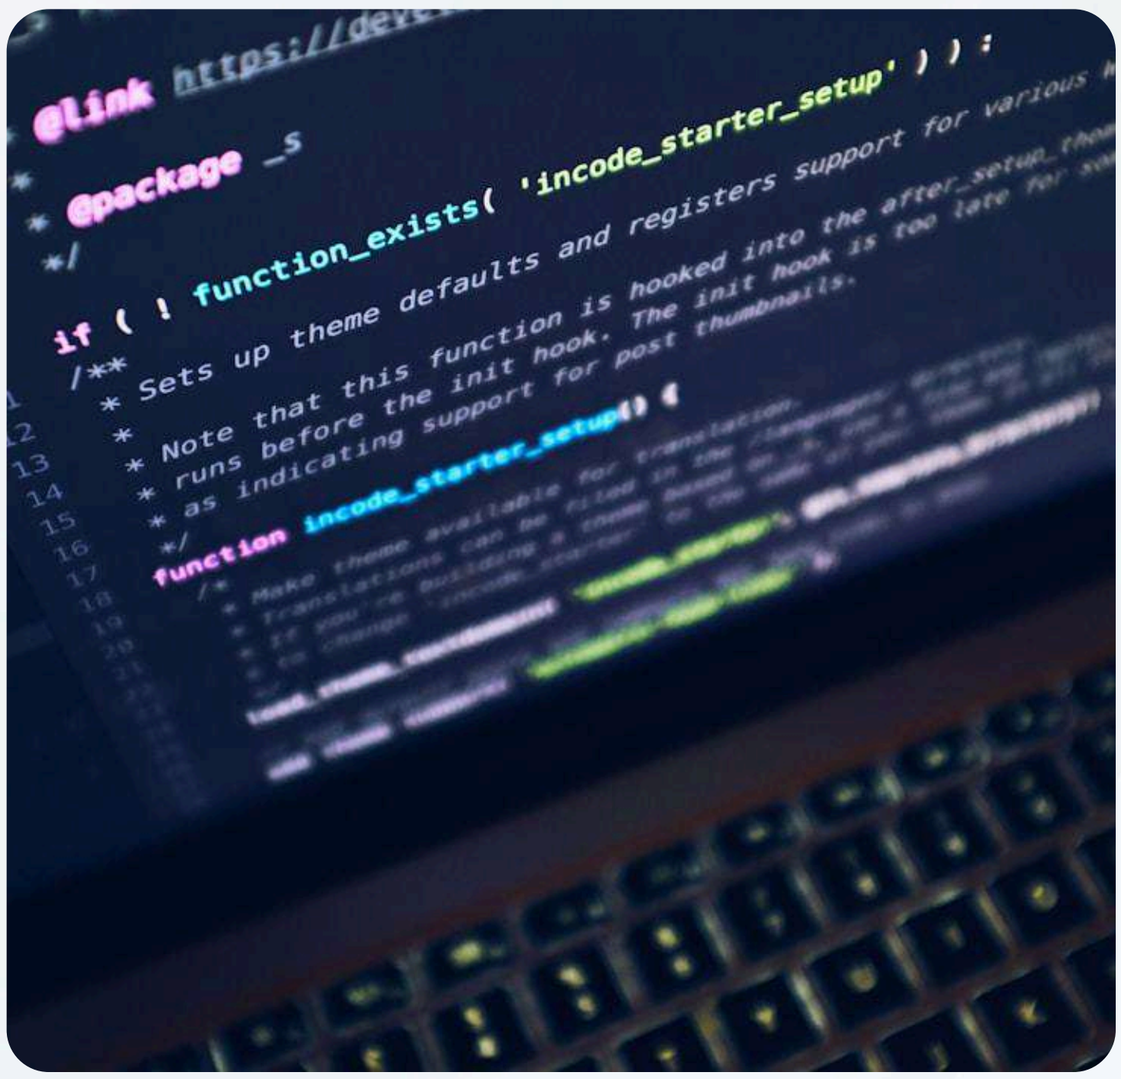
Nonprofits and educational institutions often operate with limited IT resources but complex stakeholder relationships. Mirketa's orchestration layer connects donor management systems, student information platforms, and fundraising tools, enabling AI-driven engagement, churn prediction, and outreach optimization at scale.

- 25% Higher donor retention through predictive engagement
- 40% Improvement in fundraising campaign ROI
- 50% Reduction in manual donor outreach effort
- 3x Faster grant application processing via AI workflows

Technology & SaaS

Technology companies face unique challenges: rapid product iteration, complex DevOps pipelines, and high-velocity sales cycles. Mirketa's orchestration layer accelerates every stage, from automated code review and deployment pipelines to AI-powered sales intelligence and customer success workflows.

- 35% Shorter DevOps cycle times via agentic automation
- 50% Reduction in mean time to remediate (MTTR) security issues
- 20-30% Improvement in sales pipeline conversion rates
- 40% Agent productivity boost with AI-powered case resolution



Chapter 6: The Delivery Framework

Most consulting engagements promise transformation. Few deliver it in weeks. Mirketa's delivery framework is designed around speed-to-value, with a standard timeline of 8-10 weeks from initial assessment to production launch.

Timeline	Phase	What Happens
Week 1-2	Architecture Assessment	Evaluate existing systems, data quality, and integration landscape. Identify highest-ROI orchestration opportunities for your business context.
Week 2-3	Framework Design	Design the MCP integration layer, agent routing logic, data models, security protocols, and Trust Layer governance framework tailored to your enterprise policies.
Week 3-8	Build & Connect	Develop the orchestration hub, connect APIs and legacy databases, configure multi-agent workflows, and run iterative demos with your team throughout the build cycle.
Week 8-10+	Deploy & Scale	Launch to production, train your teams, monitor agent performance in real time, and continuously optimize context windows and routing logic based on live outcomes.

Chapter 7: Mirketa vs. The Alternatives

Before selecting an AI orchestration partner or approach, enterprise leaders typically consider three paths: engaging a specialist like Mirketa, building an in-house capability, or adopting a generic AI platform. Here is how they compare across the dimensions that matter most.

Capability	Mirketa	Build In-House	Generic Platform
Time to Production	8-10 Weeks	6-18 Months	3-6 Months
Legacy System Integration	Zero Replacement	Requires Migration	Often Needs Upgrade
MCP / Open Standards	Native Integration	Custom Build Required	Vendor-Specific Only
Trust & Security Layer	Built-in, Configurable	Must Build from Scratch	Basic Governance Only
Client IP Ownership	100% Client-Owned	Client-Owned	Platform-Dependent
Industry Use Cases	Pre-Built Accelerators	Built from Zero	Generic Templates
Salesforce / Agentforce	Deep Native Integration	Requires SF Expertise	API-Level Only

The most important column is the one most organizations overlook: time to production. An in-house build that takes 6-18 months is not just slower, it means 6-18 months of delayed value, compounding opportunity cost, and the risk of building on architectural decisions that do not scale. Mirketa's pre-built accelerators and MCP integration framework consistently reduce delivery timelines by 40% compared to custom-built approaches.

Conclusion: Your Pilot Deserves to Become Production

Every enterprise AI pilot that stalls represents a missed opportunity, not just for efficiency gains, but for the competitive advantages that come from operating at the intelligence frontier of your industry.

The good news: the path from pilot to production is shorter than you think. You do not need to replace your core systems. You do not need to wait 18 months. You do not need to sacrifice security or compliance for speed.

What you need is a robust orchestration and data access layer that unifies what you already have, connecting AI agents to your existing systems through a governed, scalable architecture that your organization can trust and own outright.

> "We were convinced we'd need to replace our legacy EHR system to participate in AI. Mirketa proved us wrong. Their MCP integration layer connected our existing systems to AI agents in weeks, not the 18-month migration project we feared. We now have production AI workflows running on infrastructure we've had for years."
>
> **— Senior IT Director, Leading Healthcare Network**

Ready to Scale Your AI Initiative?

Talk to Mirketa's AI Architecture Team



mirketa
Think .Build .Transform

www.mirketa.com

1-855-Mirketa

@ info@mirketa.com

/MirketaInc /lifeatmirketa /mirketa-inc- /mirketainc173

All rights reserved. No part of this document may be reproduced, stored in a retrieval system, transmitted in any form or by any means, electronic, mechanical, photocopying, recording, or otherwise, without the express written permission from Mirketa. The information contained herein is subject to change without notice. All other trademarks mentioned herein are the property of their respective owners.