

Agentic Automation for the Enterprise

The Fourth Platform — the CIO's horizontal
AI operating system for the enterprise.



CONTENTS

01

Current State of Enterprise AI

02

Case Study: Global Top 10 Bank

03

The 80/20 of Enterprise AI

04

Introducing the Horizontal AI Operating System

05

Differentiator: Horizontal Decoupling and
The Enterprise Brain
Knowledge, Action, Governance

06

Differentiator: Vertical Decoupling and
Assembly-first Methodology
Workflows, Agents, Apps

07

Real World Example of AI OS in action

08

CIO's Choice

09

Conclusion



Rewiring Your Business to Win in the Age of Intelligence

EXECUTIVE SUMMARY

AI doesn't fail enterprises, enterprise structures fail AI

95% of Enterprise AI never reaches production. Not because the models are weak — because everything around them was built for a pre-AI world.

Most AI initiatives bolt intelligence onto fragmented systems, siloed data, hard-coded workflows, and disconnected governance. The result is predictable: pilots that demo well but never scale into repeatable enterprise value.

The fix isn't better models. It's better architecture.

80% of enterprise work today is complex, repetitive, cross-functional, and trapped between systems. Automating it takes more than chatbots, copilots, or point solutions. It takes a horizontal intelligence layer that understands enterprise context, acts across systems, and governs outcomes at scale.

That starts with horizontal decoupling: separating enterprise context from individual applications and organizing it into a shared Enterprise Brain across data, actions, and governance.

It continues with vertical decoupling: replacing code-first with assembly-first, where workflows, apps, and agents are composed once — then reused, improved, and swapped without rebuilding.

Together, they form the foundation for agentic automation: an operating system where AI moves beyond assistance and starts executing real work with context, control, and accountability.

1. MIT: The state of AI in Business 2025

2. McKinsey: The state of AI in 2025: Agents, innovation, and transformation

88%



THE
GAP



95%

of enterprises report AI use in at least one business function¹

of these AI Pilots are failing to deliver measurable ROI²



01

Current State of Enterprise AI

Enterprises are moving quickly on AI, but the underlying architecture needed to make it work in production is still missing.

How enterprises deploy AI today

- Building directly on models
- Relying on existing vendors
- Shipping point solutions in silos
- Monolithic code-first approach
- Maximizing AI usage

What breaks in production

- No clear path from pilot to production outcomes
- Intelligence stays trapped inside the app layer
- Every use case creates another version of the truth
- Logic, models, workflows, and UI are locked together
- Token costs grow without governance or ROI visibility

The core problem is that enterprises focus on the visible 20% — models, interfaces, and use cases, while neglecting the 80% required for production AI: context, governance, actionability, and architectural discipline.



02

CASE STUDY

Unified Marketing Communications & Consent at a Global Top 10 Bank

Real-life customer example **before** UnifyApps

The bank ranks in the Top 10 globally, offering a full range of financial services from retail banking (loans, cards, accounts) to wholesale banking (corporate finance, investment banking) and wealth management, serving individuals, businesses, and large corporations; internationally, known for its extensive network and digital offerings.



The bank **faced two compounding, systemic failures:** an **uncontrolled explosion** in outbound marketing communications, and a **simultaneous erosion** of reliable customer consent and preference data.



Data Context

No Enterprise Source of Truth

CURRENT STATE

- No unified, real-time view of customer consent & communication preferences.
- Consent data fragmented across 27+ systems & multiple touch-points.
- No standardized data model
- No auditable consent records.
- No real-time synchronization across banking platforms and apps.

IMPACT

- Customers **repeatedly update** preferences.
- Inconsistent & irrelevant communications.
- No reliable inputs for analytics, optimization, or AI-driven decisioning.
- Manual reconciliation required for audits and investigations.

CORE FAILURE

The enterprise lacked a **single, trusted data layer** for consent & communication—making intelligence impossible.

Action Context

No Coordinated Execution Layer

CURRENT STATE

- No centralized system to act on consent in real time across channels: 2.5B messages being sent every month—a single customer could receive over 100 messages/month.
- Communication execution fragmented across 200 disparate tools.
- No dynamic routing across channels.
- No closed-loop integration downstream.

IMPACT

- Duplicate and conflicting messages for the same customer and transaction.
- High operational cost (\$4.1M/month aggregator spend, growing 5% MoM).
- Inability to optimize delivery rates.
- Slow response to preference changes.

CORE FAILURE

The enterprise had channels—but no **action layer** to coordinate, optimize, and enforce decisions at scale.

Governance Context

No Orchestrated Governance Layer

CURRENT STATE

- No policy-driven consent enforcement across products and channels.
- No audit-ready consent trails aligned with Data Privacy and Protection rules.
- No centralized governance over 60,000+ templates.
- Weak security controls in workflows (retry limits, lockouts, session handling).

IMPACT

- Regulatory exposure and inability to demonstrate consent compliance.
- Increased risk of penalties and reputational damage.
- Manual, reactive reporting.
- Erosion of customer trust.

CORE FAILURE

The enterprise lacked a governance layer capable of enforcing policy and proving compliance—establishing trust by design.

Without **data**, the system could not understand customer intent. Without **action**, it could not respond intelligently. Without **governance**, it could not operate safely or at scale. To solve this digitally, the bank determined that it would take **12-18 months...**

03

The 80-20 of Enterprise AI

80% of the work required to bring AI into production has nothing to do with the model. The model does only 20% of the work. The other 80% is everything around the model: connecting systems, reconciling context, setting governance, building workflows, agents and apps. Today it's hand-coded for every use case and rebuilt from scratch the next. As a result, it never scales.

Moving from pilot to production requires solving that 80% once, at the enterprise level: the systems, context, orchestration, and operating foundation that make AI usable in the real world.

20%

THE MODEL

80%

EVERYTHING ELSE IT TAKES TO REACH PRODUCTION

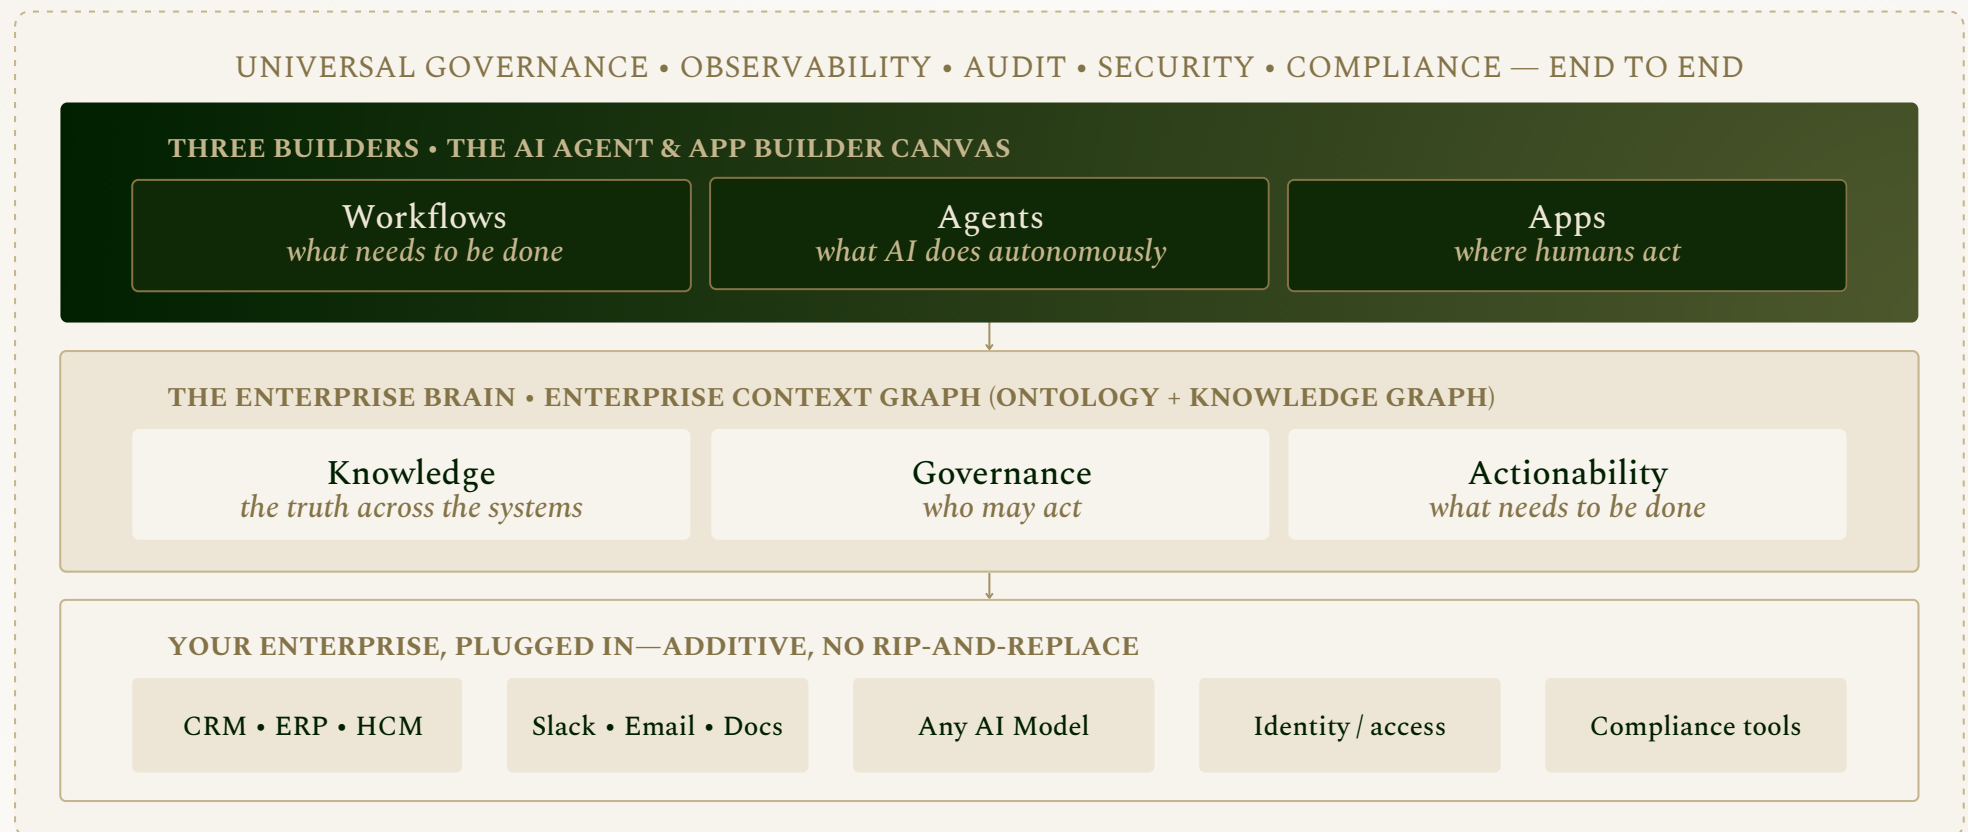


04

Introducing the Horizontal AI Operating System

Automating the 80% of enterprise work trapped between systems takes more than another AI app or agent — it takes a horizontal AI operating system where AI understands context, acts across systems, and stays governed.

At the center is the Enterprise Brain: connective tissue synthesizing knowledge, governance, and actionability, so each use case makes the next faster, cheaper, smarter. It works in layers — systems at the bottom, three contexts synthesized into the Brain, workflows, apps, and agents composed on top — powered by horizontal and vertical decoupling.



The whole platform on one page: systems in → contexts synthesized into the enterprise brain → builders on top → governance around all of it.



05

Differentiator: Horizontal Decoupling > The Enterprise Brain


Every system stores its own version of the truth — pull customer data from your CRM and ERP and they'll conflict. Most teams reconcile this by hand, for every use case: a chaotic, non-scalable multiple-versions-of-truth problem.

Horizontal decoupling fixes this at the foundation. Each use case extracts context, reconciles it, and synthesizes it into a shared Enterprise Brain. The next team doesn't start from scratch — they build on a harmonized foundation.

That context spans three layers:

- Knowledge, so AI understands the facts and relationships that matter
- Actionability, so it knows what work can be performed across systems
- Governance, so every action is controlled, compliant, and observable

You're not solving everything at once. You're ensuring every use case leaves the enterprise smarter than it found it — and that you stop renting your brain from the app and cloud giants, and start owning it.



THE PROBLEM · DUPLICATION

Context stays trapped inside each app, so every use case rebuilds the truth — and they disagree. Add more use cases and the conflicts compound, until pilots stall.

THE FIX · HORIZONTAL DECOUPLING

Context is extracted once and reconciled into the Enterprise Brain. Every new use case sits on firmer ground than the last — faster, smarter, cheaper.



The Three Contexts:

To operate in production, AI needs enterprise context decoupled from individual applications and extracted into three reusable layers across knowledge, actionability, and governance.

KNOWLEDGE CONTEXT:

Every piece of software in the Enterprise stack has access to data, but not understanding. Knowledge extraction turns scattered enterprise data into reusable enterprise context: the resources, relationships, definitions, and source-of-truth rules that define how the business works.

ACTIONABILITY CONTEXT:

AI may understand the next step, but not how to execute it across systems, fields, workflows, permissions, and dependencies. Actionability context turns process logic into reusable execution paths.

GOVERNANCE CONTEXT:

AI can recommend or act, but it does not know which policies apply, who has authority, what approvals are required, or when humans must intervene. Governance context turns controls into reusable operating logic.

Together, these contexts create a shared foundation that lets every AI reason, act, and stay governed at scale.



06

Differentiator: Vertical Decoupling > Assembly-First Methodology

Most enterprises brute-force AI with monolithic builds that bundle workflow, model, UI, and logic together. It works once, then breaks when models change, costs shift, or humans need to re-enter the loop.

Vertical decoupling separates the task from the actor. Instead of coding every solution from scratch, teams assemble workflows from atomic, reusable components — then decide step by step what AI handles autonomously and where humans stay involved.

- Swap models freely: upgrade or reroute without rebuilding the application.
- Human + AI by design: add oversight, approvals, or manual execution wherever needed.
- Compound, don't duplicate: reuse workflows, apps, agents, and logic across use cases.

This is an assembly-first methodology for deploying AI in production. The compounding effect is real: your second solution is faster than your first, your tenth costs a fraction.



The Three Builders:

Once context is decoupled, enterprises need reusable builders to turn that context into production systems: workflows, agents, and apps.

WORKFLOWS:

Defines how work moves across systems, decisions, and handoffs. Workflows turn repeatable business processes into reusable orchestration that can include both AI and humans.

AGENTS:

Executes tasks using shared context, tools, permissions, and boundaries. It makes AI workers composable, governable, and reusable across use cases.

APPS:

Provides the interaction layer for humans to review, approve, intervene, and operate. It turns interfaces into flexible experiences built on shared workflows and agents.

Together, these builders make vertical decoupling real: work is assembled from reusable components instead of rebuilt from scratch.



07

CASE STUDY

Unified Marketing Communications & Consent at a Global Top 10 Bank

Real-life customer example **after** UnifyApps

In section 02 Current Reality, we introduced this bank. As you may recall, the bank was facing a 12-18 month project to develop, test and launch their two use cases. UnifyApps enabled the bank to deliver a production-ready, regulations-compliant Communications Hub and Consent Management Portal in just two weeks.


This acceleration was possible because UnifyApps is a platform, not a project. Core capabilities—authentication, security, audit trails, and policy enforcement—were already built in, eliminating months of custom development and integration risk. Centralized cloud architecture with secure APIs enabled real-time synchronization across digital touchpoints immediately, a level of consistency that would normally require extensive, brittle, point integrations.



For CIOs, this illustrates the real value of an AI-Native operating system: **speed without shortcuts, compliance without fragility, and differentiation without risk.**



Unified Marketing Communications & Consent at a Global Top 10 Bank



08

CIO's Choice

Every CIO is now facing the same decision—whether explicitly or by default:

You can treat AI as **another tool** to be embedded into existing applications, **rented** to vendors.

Or you can treat AI as a **factory** that your leverage for products—that you **build and own**.

Or you can treat AI as a **strategic operating system**—that you partner with to **innovate and own**.

The difference between those three choices will determine whether your organization captures incremental efficiency or structural advantage.

Aspect	Renting Intelligence	Build Intelligence on Your Own	Build Intelligence on with an AI OS Partner
WHAT You Get	Vendor copilots embedded in their apps	Your own factory with custom agents	Your own Enterprise OS with custom agents
WHERE Intelligence Lives	Trapped inside vendor systems	Dispersed In each App and Agent	In your centralized Knowledge Graph
WHO Controls Context	The vendor	You	You
WHO Controls Improvement	The vendor's roadmap	Your teams, on an extended timeline	Your teams, on an accelerated timeline
SWITCHING COST	High—intelligence doesn't transfer	Medium—intelligence difficult to transfer	Low—intelligence is portable
CUSTOMIZATION	Limited to vendor options	Limited to your resources	Unlimited—build exactly what you need
CROSS-SYSTEM INTELLIGENCE	Impossible—silos don't talk	Difficult as silos still exist	Native—the whole point of the architecture

The Verdict:
The Architecture of
Sovereignty

“In the Copilot Era, we bought software that had logic trapped inside it. In the AI OS Era, we extract that logic. We decouple the Context, Action, and Conscience from the application layer, creating a **single, sovereign intelligence that we own and innovate, not rent.**”



09

Conclusion

Every major era of enterprise computing has belonged to the CIO who recognized the inflection point early.

The shift to client-server.

The move to the cloud.

The rise of digital platforms.

AI is the next inflection...

This is not another application cycle. It is a **structural change** in how work, decisions, and scale operate. The constraint is no longer compute or data. It is cognition.

The Choice

CIOs now face a clear choice.

You can continue optimizing the past—adding AI features inside siloed systems, increasing integration complexity, and accepting rising cost with diminishing returns.

Or you can partner to build the next foundation—an AI operating layer that unifies data, action, and governance, and allows intelligence to scale safely across the enterprise.

One path preserves control in the short term. The other creates **leverage for the next decade**.

The Mandate

AI does not remove the need for leadership. It concentrates it.

When execution becomes cheaper, direction becomes the scarce asset. When agents can act, governance must be architectural. When intelligence is abundant, intent defines outcomes.

This is now a CIO responsibility—not as a technology buyer, but as the architect of enterprise intelligence.

The Call to Action

Rather than starting with a transformation roadmap, start with one governed use case.

Choose one workflow that matters. Centralize its context, enable safe action, enforce policy by design, and prove the pattern. Then scale it horizontally.

This is how AI moves **from pilots to production**.

This is how **cost decouples from headcount**.

This is how **enterprises become AI-Native**—by design, not accident.

The AI Era has already started. The only question left is whether your operating system is ready.

Now is the moment to decide—and to lead.



You can continue optimizing the past—

Or you can build the next foundation and allow intelligence to scale safely across the enterprise.



CEO's Note



Ragy Thomas
Chairman, CEO & Co-founder, UnifyApps

“

The shift from a digital-native enterprise to an AI-Native enterprise is happening now. Will you lead it?



Let's make your
Business AI-Native.

