DOCUMENT

The Case for Agentic Technology Delivery.

Why the next structural advantage in enterprise technology isn't a better tool — it's a fundamentally different model for building software, creating AI-native products, accelerating legacy systems, and automating operations.

THE CORE THESIS

Software delivery fails not because teams lack talent, but because the process wastes most of it. 73% of enterprise software projects overrun budget or timeline. The usual explanation is "bad estimates" or "scope creep." The real cause is structural: ambiguous requirements pass unchallenged into engineering, architecture decisions happen too late, testing gets compressed, and defects surface where they're most expensive to fix.

The same pattern plays out in process automation. RPA automated the steps — but never understood the goal. When exceptions arise, bots stop. When processes change, scripts break.

And when enterprises try to build AI-native products, they find that calling an LLM API is straightforward — but designing multi-agent systems with domain intelligence, enterprise governance, and production reliability is an entirely different engineering challenge.

Predikly has spent a decade inside these problems. That experience is now the foundation for four interconnected disciplines — each powered by intelligent agents that own phases of work with defined authority, measurable outputs, and human oversight.

AUDIENCE	CLASSIFICATION	READING TIME	FORMAT
CTO · CPO · CFO · CIO	Confidential	12–15 minutes	PDF · 8 pages

FOUR DISCIPLINES

One Company. Four Ways to Deploy Agentic Intelligence.

Each discipline addresses a different enterprise challenge. All four share the same architectural principles, governance framework, and delivery model.

Agentic Solutions

Build production-grade AI-native products where multi-agent architectures own domain intelligence. Orchestrated agent hierarchies, domain ontologies, ML model integration, and enterprise governance — from concept to production.

Agentic Product Development

Nine intelligent agents across the full SDLC — from requirements through production. Each agent owns a delivery phase with defined authority, measurable outputs, and human oversight. The way software should be built.

Enhancement Acceleration

AI-powered legacy code analysis, enhancement acceleration, and intelligent support on codebases already in production. Understand what exists deeply, then compress every phase of enhancement delivery.

Agentic Process Automation

Beyond RPA — goal-oriented agents that understand the purpose of a process, handle exceptions within boundaries, make bounded decisions, and improve with every execution cycle.

17+	9	↑ 2.5×	↓ 80%
Specialized agents in production	SDLC agents per engagement	Enhancement delivery speed	Exception escalations

AGENTIC SOLUTIONS

Build AI-Native Agentic Solutions with Agent Architectures.

Production-grade platforms where multi-agent architectures own domain intelligence — not chatbots bolted onto existing software.

Four-Layer Architecture Framework

L4 Experience Layer	**API Gateway + Frontend + MCP** User-facing interfaces, chat UIs, planning wizards, dashboards, and MCP servers for external AI integration. FastAPI · Nuxt / Next.js · MCP · SSE
L3 Intelligence Stack	**Agent Orchestration + ML Models + RAG** Multi-agent hierarchies, structured tool registries, Bayesian models, vector search, and forecasting — the decision engine. Claude API · GPT-4 · PyMC · XGBoost · GraphRAG
L2 Domain Model	**Typed Semantic Ontology** Object types, link types, taxonomies, and business rules. The layer that makes AI domain-specialized, not general-purpose. Object registry · Taxonomy · Hierarchy engine · Rules
L1 Data Foundation	**Enterprise Data Platform — Single Source of Truth** ACID transactions, governance, compute, and model registry. Platform-flexible across enterprise data infrastructure. Databricks · Snowflake · BigQuery · PostgreSQL

What We Build

Multi-Agent Orchestration

Hierarchical agent architectures — orchestrators routing to domain specialists, specialists invoking workers.

Agent holarchies · Intent routing

Domain Ontology Design

Typed semantic models encoding vocabulary, relationships, and rules. Domain-specialized intelligence.

Object types · Business rules

ML Model Integration

Bayesian inference, gradient boosting, deep learning, forecasting — integrated into agent decision loops.

PyMC · XGBoost · PyTorch

RAG & MCP Integration

Vector search, GraphRAG for grounded responses. MCP servers for external AI system interoperability.

Vector search · MCP protocol

AGENTIC PRODUCT DEVELOPMENT

Nine Agents. Six Phases. One Continuous Loop.

Each agent has a defined domain, measurable outputs, and explicit authority boundaries. They don't replace your engineers — they take ownership of the systematic work.

01 DISCOVER

Requirements structured, gaps surfaced, scope validated

Stakeholder input transformed into traceable, structured requirements. Contradictions and scope gaps identified before engineering begins.

ClariX — BA Agent

↓ 70%
Requirements phase duration

02 DESIGN

Architecture decided with evidence, not opinion

Multiple options with documented trade-offs. Tech stack evaluated against team capability and long-term cost.

StratiX — Architect · TrackX — Tech Lead

↓ 70%
Architecture decision time

03 BUILD

Frontend, backend, and data layer generated in parallel

Production-ready components, secure APIs, and versioned schemas. Data contracts defined before integration.

CraftX · ForgeX · VaultX

↑ 3×
Component delivery speed

04 VALIDATE

Quality moved upstream — defects caught at source

Tests before code. Self-maintaining regression suites. Defect prediction identifies highest-risk changes.

VeriX — QA Agent

↓ 65%
Defects in production

05 SHIP

Deployed with confidence, monitored with context

CI/CD pipelines optimised. Deployment risk assessed. Incidents triaged with full codebase context.

FleetX — DevOps

↓ 70%
Mean time to resolution

GuardX — Security Agent (Cross-cutting, all phases)

Threat modelling from requirements. Real-time vulnerability scanning. Compliance validation against SOC2, HIPAA, PCI-DSS, GDPR. Dependency auditing on every build.

Threat modelling · Real-time scanning · Compliance · Dependency audit

THE ECONOMICS

Where Does the Money Actually Go?

Most engineering spend doesn't go to engineering. It goes to the overhead between engineering — the waiting, the rework, the miscommunication, the context-switching.

BEFORE	WITH AGENTIC DELIVERY
3–4 week discovery cycles. Ambiguous stories enter sprint planning. Rework surfaces in week 3.	48 hours to structured, validated requirements. Gaps surfaced before a line of code is written.
Weeks of architecture debate. Single option presented as inevitable. Trade-offs undocumented.	3 options with documented trade-offs per decision. ADRs generated automatically.
Frontend/backend misalignment. Schema surprises. Security bolted on at review.	Parallel generation with data contracts defined upfront. Security from line one.
Test cases written after code. Regression suites decay. Production is the real test environment.	Tests before code. Self-maintaining regression suites. Defects caught at source.

↓ 58%
Overall delivery cycle reduction

↓ 65%
Defects reaching production

↑ 3.2×
Team output per sprint

↑ 91%
On-time delivery rate

Metrics measured across active engagements. Individual outcomes depend on team size, codebase maturity, and process complexity.

ENHANCEMENT ACCELERATION

Codebases That Already Exist. Delivery That Doesn't Slow Down.

AI-powered legacy code analysis, enhancement acceleration, and intelligent support — using Claude Code, OpenAI Codex, and Gemini Code Assist alongside our nine-agent ecosystem.

Legacy Code Analysis

AI-powered architecture mapping that turns months of archaeology into days.

↓ 60% Time to map architecture

Enhancement Acceleration

Agents and frontier coding models compressing delivery cycles on existing systems.

↑ 2.5× Enhancement delivery speed

Intelligent Support

AI-powered triage using actual codebase context — not pattern-matching on symptoms.

↓ 45% Mean time to resolution

AGENTIC PROCESS AUTOMATION

Beyond RPA. Agents That Understand Goals.

Intelligent agents that own end-to-end processes — handling exceptions, making bounded decisions, and improving with every execution cycle.

Intelligent Workflow Agents

Agents that own entire processes — understanding goals, adapting when conditions change, routing work intelligently based on context.

Exception Handling Agents

The 20% of cases that break every RPA implementation — assessed against policy and context, resolved within defined authority.

Document Intelligence

Agents that read, understand, extract, validate, and route document content with contextual comprehension beyond OCR.

↓ 80%	↑ 5×	↓ 60%	↑ 94%
Exception escalations	Process coverage vs RPA	Total cost of automation ownership	Straight-through processing rate

GOVERNANCE & EVIDENCE

The Question Isn't Whether AI Can Help. It's Whether You Can Trust It.

Every enterprise deploying AI agents faces the same governance challenge. We built for that question first — not as an afterthought.

Every output scored

Agent outputs evaluated against correctness, completeness, and project alignment. Continuously — not as a one-time setup.

Drift detected early

Behaviour monitored against baselines. When models update or context shifts, alerts fire before delivery is affected.

Agents tested like software

Defined test cases, expected outputs, regression suites. No prompt or model change reaches production untested.

Humans stay in control

Confidence thresholds, mandatory review gates, explicit escalation triggers. Agents know when to hand off.

Every decision traceable

Full input-output logging with decision rationale. Exportable audit reports. RBAC and data residency compliance.

Full version control

Every prompt, model version, and configuration change tracked. Instant rollback if a version introduces regressions.

WHAT THIS MEANS FOR YOUR LEADERSHIP TEAM

Four conversations. Four different priorities.

CIO / COO

Governance is in the architecture, not a retrofit.

Every agent action logged with full traceability. Decision boundaries are explicit. Your risk and compliance teams won't need to add controls after the fact.

CFO

Your forecasts start reflecting reality.

Agents enforce structure at every phase — validated requirements, tracked dependencies, quality gates. Predictability becomes structural, not promised.

CTO

Your best people work on problems that matter.

Senior engineers stop drowning in scaffolding and process overhead. Agents own the systematic work. Technical leaders focus on the decisions that determine whether the product holds.

CPO

Smaller bets. Faster learning. Lower cost of being wrong.

When build cycles compress from quarters to weeks, you stop placing large bets on static specs and start running experiments against real user behaviour.

EVIDENCE

AGENTIC SOLUTIONS · LIVE

MySavi.ai

Greenfield AI-native marketing intelligence platform. 17+ specialized agents in a 5-layer holarchy. Bayesian MMM, GraphRAG, enterprise RBAC. Predikly is the development partner from line one — concept to production.

ENHANCEMENT · ACTIVE

Enterprise Legacy Acceleration

2× enhancement delivery speed on decade-old codebases. Using Claude Code, OpenAI Codex, and Gemini Code Assist alongside our nine-agent ecosystem. Full architectural visibility within weeks.

FOUR DISCIPLINES

Agentic Solutions
Build AI-native products with multi-agent architectures, domain ontologies, and enterprise governance.

Product Development
Greenfield builds with agentic AI from the first commit. Nine agents across the full SDLC.

Enhancement Acceleration
Understand, enhance, and accelerate existing codebases using frontier coding models.

Process Automation
Beyond RPA — goal-oriented agents for enterprise operations.

See This Applied to a Real Challenge.

The best way to evaluate whether agentic delivery changes anything for your organisation is to see it applied to something specific — your codebase, your team structure, your delivery timeline, or your product vision.

Book a 30-Minute CXO Briefing →

Or email us directly at prediklyleadership@predikly.com

WHAT HAPPENS
A direct conversation about your current challenges and where agentic agents would have the highest impact.

WHO YOU SPEAK WITH
Senior technical leadership at Predikly. No intermediaries. No pitch decks in the first meeting.

WHAT YOU LEAVE WITH
A clear assessment of fit, recommended engagement model, and projected timeline — regardless of whether we work together.