

Executive Guide 2026

PRODIGY

Ciklum's AI Platform, AI Accelerators and Delivery Methodology

Built for Enterprise and Tech Leaders



Table of Contents:

The challenge

Why is AI Transformation Failing Today? ----- 03

Introducing PRODIGY

Introducing PRODIGY ----- 04

Vision: AI Ingenuity. Enterprise Reimagined. ----- 05

Why PRODIGY Matters for Enterprise AI ----- 06

Why PRODIGY Wins Where Others Stall ----- 07

Platform and accelerators

Core Components of PRODIGY ----- 09

Accelerator Ecosystem ----- 09

Technical Accelerators ----- 11

Horizontal Accelerators ----- 13

Vertical Accelerators ----- 13

Delivery and governance

Delivery Methods ----- 14

PRODIGY PDLC ----- 15

AgentKit ----- 16

AI Product Incubator ----- 17

Strategy Clinics ----- 17

Platform Architecture ----- 19

Security and Governance ----- 19

Value and outcomes

The Enterprise Value of PRODIGY ----- 20

What Makes PRODIGY Different ----- 21

Enterprise Use Cases ----- 23

Proof in Practice: AI-First Healthcare Product Engineering ----- 24

Measuring AI Success with PRODIGY ----- 25

How Organisations Can Start ----- 26

Summary ----- 27

Why is AI Transformation Failing Today?

Enterprise leaders today have moved on from asking ‘Does AI have potential?’ They are asking why that potential is still so hard to turn into repeatable business value.

 **AI initiatives don’t begin with validated problems.**

Teams move quickly into pilots, but if the problem has not been tested against customer insight, operational reality, data readiness, and commercial impact, the solution can look impressive while solving the wrong thing.

 **Pilots are easy to start. Production systems are harder to scale.**

A proof of concept can work in a controlled setting. Production AI has to work across live systems, fragmented data, security requirements, governance controls, user workflows, and changing model ecosystems.

 **Every AI initiative starts from scratch.**

Teams rebuild the same foundations across data integration, knowledge grounding, security, testing, and deployment. This slows delivery and makes AI harder to scale beyond isolated use cases.

 **The traditional buy-or-build choice creates a trade-off.**

Off-the-shelf AI can be faster, but it can limit flexibility, control, and differentiation. Building from scratch gives more ownership, but it usually takes longer, costs more, and requires scarce technical expertise.

 **The right signals often arrive too late.**

Enterprises may already have the data they need, but it is often spread across different systems and difficult to act on quickly. By the time risks or opportunities become visible, the best moment to respond may already have passed.

That is the gap PRODIGY is built to close.

Introducing PRODIGY

Ciklum's PRODIGY combines an enterprise AI platform, accelerator ecosystem, and delivery methodology in one, designed to help organisations build and operate AI-powered products and systems at scale.

It consists of three key elements:



A structured AI-native delivery method



A modular ecosystem of AI accelerators



A platform architecture for orchestrating agents, knowledge, and enterprise systems

Together, these components form an operating model that allows organisations to move from isolated AI experiments to systematic, enterprise-wide AI adoption.

PRODIGY supports the full lifecycle of AI-enabled products and systems:



Strategy and AI opportunity identification



Product and experience design



Software engineering and AI model integration



QA, deployment, and operations



Continuous optimisation and value operations

In practice, PRODIGY is designed to move business metrics that matter, including delivery speed, engineering throughput, operational cost, time to production, and product scalability.

In live engagements, PRODIGY-enabled approaches have supported outcomes such as 40% faster product development lifecycle delivery, 25% operational cost reduction, and more than 60% faster migration work against baseline delivery.

Vision: AI Ingenuity. Enterprise Reimagined

The vision behind PRODIGY is to help organisations transition from digital-first enterprises to AI-first enterprises.

Organisations today apply AI in fragmented ways. Individual teams experiment with models, deploy chatbots, or automate narrow workflows. While these initiatives can produce value, they rarely transform how organisations design products or operate their platforms. PRODIGY addresses this challenge by introducing a coherent framework for enterprise AI delivery.

Instead of isolated solutions, organisations adopt a structured model that embeds AI across the entire product lifecycle and enterprise architecture. That means:

-  AI embedded into every stage of product development
-  AI agents augmenting engineering and operational workflows
-  Enterprise knowledge systems that ground AI decisions in a real business context
-  Governance and security integrated from the start

GenAI Engagements

>30
GenAI projects

>30
In-production GenAI builds

>33
Industry Recognitions within the last 18 months

>150
Powerful computational servers in 3 jurisdictions

>350
Use cases identified

>19
Partnerships



Why PRODIGY Matters for Enterprise AI

AI-enabled systems are becoming more complex to design, build, govern, and scale.

Traditional digital engineering approaches often require large portions of the solution to be built from scratch. This model becomes less efficient when AI initiatives need to integrate across models, infrastructure, governance frameworks, and enterprise data platforms.

PRODIGY introduces reusable platform components that allow organisations to start from a pre-built foundation rather than a blank slate.

This helps enterprise teams:

-  Accelerate delivery timelines
-  Increase consistency across projects
-  Reduce engineering overhead
-  Support higher-value AI transformation initiatives

For enterprise leaders, PRODIGY supports four core priorities:



Moving from experimentation to execution

By combining consulting, accelerators, and platform capabilities, PRODIGY helps organisations move beyond isolated AI pilots and toward AI transformation programmes.



Building reusable capability

PRODIGY formalises proven approaches into reusable assets, including:

Accelerators Delivery frameworks Architecture blueprints Governance models



Enabling scalable AI delivery

AI systems introduce new complexity across data, models, orchestration, and governance. PRODIGY provides the structured architecture required to manage this complexity.



Creating differentiation through speed and control

Many organisations have access to the same AI models and cloud platforms. The difference lies in how quickly and consistently they can turn those capabilities into working products and systems.

Why PRODIGY Wins Where Others Stall

Enterprise AI initiatives often stall for the same reasons. The proof of concept works, but the data is not ready. The model performs in a controlled environment, but cannot be governed in production. The business sees potential, but cannot connect activity to measurable value.

PRODIGY is built around the capabilities required to get past those barriers.

Why Ciklum?

This isn't Ciklum's first AI cycle. Ciklum has run a dedicated AI R&D function since 2015, became NVIDIA's first European partner, delivered one of the first production-grade generative AI solutions in Germany, and built one of the first agentic solutions in production back in 2023. PRODIGY is the product of that track record throughout the years.

Why PRODIGY?

Every major competitor already has access to the same models and cloud platforms. What's scarce is the ability to turn that access into secure, production-ready systems repeatedly. All without rebuilding the foundation every time. PRODIGY exists to close that specific gap.

Why now?

Enterprise AI is moving faster than typical software cycles can absorb. Models and vendor preference are shifting within months, not years, and cost and context management have become bigger constraints on scale than access to AI itself. Those who wait for the landscape to settle risk building on assumptions that are already out of date.



Why Ciklum?


-  Dedicated AI R&D since 2015
-  NVIDIA's first European partner
-  Production-grade GenAI delivered in Germany
-  Agentic AI in production since 2023

Why PRODIGY?

-  Reusable accelerators
-  AI-native delivery methods
-  Platform architecture for governed scale
-  Built for repeatable enterprise delivery

Why Now?

-  AI cycles are moving faster
-  Model and vendor choices keep shifting
-  Cost and context now define scale
-  Waiting increases delivery risk



Core Components of PRODIGY

Accelerator Ecosystem


Accelerators are pre-built capabilities that reduce the effort required to implement AI solutions.

They are modular and composable, meaning organisations do not need to adopt the full platform at once. Instead, accelerators can be selected and integrated depending on the business problem, technical environment, and level of AI maturity.

PRODIGY consists of three primary layers:

Layer	Purpose
Accelerators	Reusable technical and domain capabilities
Methods	Structured delivery frameworks
Platform	Architecture for orchestration and integration

PRODIGY accelerators are grouped into three areas:

-  **Technical accelerators** that support the foundations of AI delivery
-  **Horizontal accelerators** that apply across common enterprise workflows
-  **Vertical accelerators** that address industry-specific problems

Accelerators help organisations avoid rebuilding the same foundations for every AI initiative, while still allowing solutions to be adapted to their own systems, governance requirements, and business priorities.

Ciklum's AI Platform: PRODIGY

The **AI platform and methodology** for delivering successful AI:



Drives both AI SDLC and the creation of AI solutions



Accelerates delivery across the entire lifecycle



Platform of accelerators (Ciklum IP) provide clear market differentiators




Experience first approach to creating AI applications

Prodigy Methodology



AI Ideation



AI Experience



AI Delivery



AI Infrastructure

Prodigy Platform

Ciklum IP:

Vertical Accelerators

(e.g Account re-engager)

Horizontal Accelerators

(e.g. AI Support Agents)

Technical Accelerators

(e.g. AI Knowledge Engine)

Ciklum Proprietary Knowledge:



Domain Knowledge



Prompting Methods



Agent Optimisation Methods



Security

Technical Accelerators

Technical accelerators form the foundational layer of PRODIGY. They support the orchestration, integration, knowledge grounding, and delivery workflows required to build enterprise AI systems at scale.



Agentic Framework

The Agentic Framework enables AI agents to participate directly in the software development lifecycle.

Agents can support activities such as requirements analysis, backlog creation, code generation, unit test production, QA scenario validation, and release readiness checks.

By embedding AI into delivery workflows, the framework helps teams move toward a more AI-native SDLC, improving speed and consistency while keeping human judgement in control.



Knowledge Engine

The PRODIGY Knowledge Engine helps ground AI in a real enterprise context.

It connects to project documentation, tickets, codebases, backlogs, transcripts, third-party documents, and regulatory information, giving teams a more reliable way to use AI across discovery, requirements, delivery, testing, and handover.

For enterprise leaders, this addresses one of the biggest barriers to AI adoption: generic AI outputs that are not grounded in how the organisation actually works.

The Knowledge Engine supports faster onboarding, automated gap analysis, ticket creation, documentation generation, traceable outputs, stronger requirements, and integration with existing tools such as Jira, Confluence, GitHub, Drive, and test platforms.

With the Knowledge Engine, AI becomes more useful in complex enterprise environments where context, traceability, and accuracy matter.



Integration and developer tooling

PRODIGY also includes technical accelerators that simplify integration with enterprise systems, AI agents, and model-context ecosystems.


These capabilities help teams connect AI to existing tools and workflows faster, while maintaining consistency across development, orchestration, and deployment.

PRODIGY Knowledge Engine

Grounding AI in your enterprise context for accurate, relevant, and traceable outcomes.

Enterprise Context Inputs

- **Project Documentation:** PRDs, specs, designs
- **Tickets & Backlogs:** Jira, Azure DevOps, etc.
- **Codebases:** Repos, libraries, frameworks
- **Transcripts:** Meetings, calls, reviews
- **Regulatory Information:** Policies, standards, laws
- **Internal Policies:** Security, compliance, HR
- **Third-party Documents:** Whitepapers, contracts, reports
- **Enterprise Systems:** Jira, Confluence, GitHub, Drive, etc.



Intelligent Outputs

- **Traceable AI Outputs:** Citations, sources, auditability
- **Stronger Requirements:** Clear, complete, unambiguous
- **Faster Onboarding:** Context-rich, role-specific
- **Automated Gap Analysis:** Identify risks, inconsistencies
- **Better Documentation:** Auto-generated, always current
- **Smarter Decisions:** Insights grounded in reality

Horizontal Accelerators

Horizontal accelerators provide reusable AI capabilities that apply across industries and business functions.

They support common enterprise needs such as document processing, customer support, IT operations, workflow automation, analytics, and decision intelligence.

Rather than creating disconnected point solutions, these accelerators help organisations embed AI into repeatable, high-volume workflows where speed, consistency, and control matter.

Examples include:



Intelligent document processing



AI-powered customer experience



Process automation



Data and analytics acceleration



IT and operational support agents

The purpose of horizontal accelerators is to help enterprise teams improve productivity and decision-making without adding another fragmented tool layer.

Vertical Accelerators

Vertical accelerators provide industry-specific AI capabilities designed around the needs of enterprise sectors where Ciklum has deep delivery experience.

Industries supported include:



Banking & Financial Services



Retail & Consumer Goods



Automotive & Manufacturing



Hi-Tech



Travel & Hospitality



Healthcare & Life Sciences



Vertical Accelerators continued...

In Banking & Financial Services, churn costs UK and US banks an estimated \$350 billion a year, and acquisition typically costs roughly five times more than retention. PRODIGY's vertical accelerators are built to catch disengagement signals early, before that cost is locked in.

In Healthcare & Life Sciences, performance against CMS Star Ratings – which govern over \$590 billion in Medicare Advantage bonus payments – can swing per-patient margin by 20–30%. PRODIGY helps plans forecast and act on these measures before the annual measurement window closes, rather than after.

Other typical use cases include fraud detection, patient service automation, supply chain optimisation, and personalised retail experiences.

Delivery Methods

Alongside accelerators, PRODIGY includes structured delivery methods that guide organisations through AI transformation.

These methods ensure that AI initiatives move systematically from idea to production.

Many enterprise AI programmes fail at the handover between experimentation and scale. A proof of concept may work in a controlled setting, but production requires a different level of readiness across systems, data, compliance, operations, and user adoption.

PRODIGY helps structure that journey.

It gives teams a repeatable approach for identifying opportunities, designing human-AI experiences, building AI-enabled systems, deploying safely, and continuously improving after launch.



PRODIGY PDLC

The PRODIGY Product Development Lifecycle (PDL) is the core delivery method within the platform.

It unifies a structured sequence of delivery steps covering the entire lifecycle of AI-enabled products.

The lifecycle includes six key phases:

01



AI Ideation

Identification of opportunities through structured ideation workshops.

02



Strategic Diagnostics

Evaluation of organisational readiness, business impact, and feasibility.

03



Discovery & Prioritisation

Use-case prioritisation using frameworks such as RISE.

04



AI Experience Design

Design of human-AI interactions and intelligent workflows.

05



AI-Native SDLC

Design of human-AI interactions and intelligent workflows.

06



Delivery & Support

Deployment, monitoring, and continuous improvement.

AgentKit


AgentKit is an enterprise platform within PRODIGY for deploying and governing AI agents directly within an organisation's own infrastructure. Rather than relying on cloud-hosted agent services, AgentKit runs entirely within a private environment, so that no data leaves the organisation's perimeter during agent operations.

The platform is built around four core capabilities:



Private Agent Cloud:

Agents run entirely within your own environment with no external dependencies.



Agent Orchestration Layer:

Supports multi-step workflows, conditional logic, and configurable human approval steps.



Enterprise Control Plane:

Centralised governance, role-based access controls, single sign-on, audit trails, and real-time visibility across all agent activity.



Development Acceleration:

Pre-built agent templates, natural language workflow creation, and instant updates without redeployment.

Workflows are defined in plain language and automatically converted into structured, branching processes that can connect to enterprise systems such as Salesforce, Jira, and any MCP-compatible tool.

AgentKit aligns with SOC 2, ISO 27001, HIPAA, and GDPR, making it suitable for regulated industries. Typical applications include support automation, release checks, internal knowledge support, and governed workflow automation.

AI Product Incubator

The AI Product Incubator is designed to rapidly validate and launch new AI products.

It focuses on accelerating the early stages of innovation by enabling teams to:



Build prototypes quickly



Validate market opportunities



Test user needs



Launch MVPs faster



Move promising ideas into production with less delivery risk

Rather than spending months on abstract strategy or disconnected experimentation, the incubator creates a structured path from concept to validated AI product opportunity.

In one banking engagement, the incubator process moved from problem discovery to two working, user-tested prototypes in 24 hours.

Strategy Clinics

Strategy Clinics are executive workshops designed to align organisations on their AI strategy.

These sessions typically focus on:



Defining an AI vision



Identifying priority use cases



Selecting platform architecture



Assessing readiness and risks




Building a transformation roadmap

They often serve as the starting point for PRODIGY engagements.

The clinics help create clarity around where AI can deliver value, what needs to be prioritised, and how the organisation can move from experimentation to scalable execution. Particularly valuable when multiple AI initiatives are already underway, but the organisation lacks a clear enterprise-level view of what should be funded next.

PRODIGY Platform Architecture

A governed architecture for connecting agents, knowledge, models, and enterprise systems



Governance Across the Stack

- Identity
- RBAC
- Audit trails
- Security
- Compliance
- Observability

PRODIGY gives enterprises a governed way to connect users, agents, knowledge, models, and existing systems.



Platform Architecture

Underneath the accelerators and delivery methods sits a platform architecture built to support large-scale AI delivery across enterprise environments. Key architectural layers include:



Presentation interfaces for users and developers



Agent orchestration and workflow management



Knowledge and intelligence layers



Infrastructure and integration systems

These layers communicate through governed APIs and shared observability frameworks — enabling secure, scalable deployment of AI capability rather than a patchwork of disconnected tools.

Security and Governance

Enterprise AI introduces new security risks, such as prompt injection, data leakage, and uncontrolled model behaviour.

PRODIGY addresses these challenges through built-in governance mechanisms.

Key capabilities include:



Enterprise identity and access integration



Full auditability of agent interactions



Execution environment isolation



Policy-based governance controls

The platform is designed to support compliance frameworks including SOC 2, ISO 27001, HIPAA, and GDPR.



The Enterprise Value of PRODIGY

PRODIGY delivers value across the areas that matter most when enterprise AI moves from experimentation into production: speed, productivity, quality, risk, and flexibility.

Faster time to value

Pre-built accelerators reduce the time required to move from idea to working solution. Instead of starting every AI initiative from a blank slate, teams can build from reusable components and architecture foundations.

Increased productivity

AI-augmented workflows improve efficiency across research, design, engineering, QA, support, and operations, helping teams reduce repetitive work while maintaining human control over quality and judgement.

Higher-quality outputs

By connecting AI systems to enterprise knowledge, documentation, codebases, and approved data sources, organisations can reduce the risk of generic or inaccurate responses.

Reduced delivery risk

Governance, monitoring, security controls, and architecture patterns help organisations move AI into production more safely, especially when systems interact with sensitive data, regulated workflows, or business-critical platforms.

Greater flexibility

PRODIGY is designed to work across models, providers, platforms, and enterprise systems, helping organisations avoid unnecessary vendor lock-in.

What Makes PRODIGY Different

Enterprise leaders often face a familiar AI decision: buy an off-the-shelf solution or build from scratch.

Buying can be faster, but it often limits flexibility, differentiation, and control. Building can offer ownership, but it usually requires more time, investment, and technical effort.

PRODIGY creates a third route.

It allows organisations to build with reusable accelerators, delivery methods, and platform components while still shaping the solution around their own systems, workflows, governance needs, and business priorities.

This helps organisations:

-  Shorten implementation timelines
-  Reduce delivery risk
-  Maintain control over architecture and roadmap
-  Scale from focused use cases to broader transformation

PRODIGY is also vendor-agnostic by design, supporting thousands of models and more than 100 providers, allowing organisations to adopt advanced AI technologies while remaining open, modular, and future-ready.

PRODIGY is changing Buy vs. Build for clients

	 Buy	VS  Build	VS  Build with PRODIGY
 Implementation Speed	 Rapid implementation time	 Slow and Complex Implementation	 Shortened implementation through Accelerators
 Cost	 High ongoing SaaS Costs	 Large time and resource investment	 Shortened implementation through Accelerators
 Flexibility	 Locked into Supplier roadmap	 Personalised Roadmap	 Personalised Roadmap
 Differentiative Potential	 Same capability as other customers	 Opportunity to build differentiative IP	 Opportunity to build differentiative IP

Enterprise Use Cases

PRODIGY supports high-value AI use cases where speed, governance, integration, and business impact all matter.



Industry-specific AI solutions

Vertical accelerators support sector-specific challenges, such as customer engagement in retail banking or performance optimisation in healthcare.




Enterprise workflow automation

Horizontal accelerators apply across business functions, including customer support, IT operations, sales enablement, document processing, and finance operations.



AI-native software delivery

The platform also supports the software delivery lifecycle itself, helping engineering teams use AI to accelerate requirements analysis, backlog creation, code generation, testing, QA validation, and release readiness.



Proof in Practice: AI-First Healthcare Product Engineering

A US healthcare technology organisation shows how PRODIGY can support AI-first product engineering in a complex healthcare environment.

The challenge was to build a privacy-compliant platform for personalised targeting across millions of consumers, with an intuitive interface for marketers to engage with customer profiles and behaviours.

Ciklum supported the client through AI-driven product engineering and AI-first PDLC enablement, using capabilities such as AI-augmented requirements processing, agentic retrieval-augmented generation, AI-assisted planning and estimation, and mock data and test scenario generation.

The engagement delivered:



Key Takeaway

PRODIGY can help accelerate complex, regulated product delivery while maintaining the structure needed for privacy, scale, and business-critical performance.

Measuring AI Success with PRODIGY

For enterprise leaders, AI investment only matters if it can be measured. PRODIGY is built to move the metrics that matter most:

These are the numbers PRODIGY is designed to move, and this is what enterprise leaders should expect their teams to report back on.



Delivery speed:

Engagements typically see 30% faster delivery of new features, rising to 60–80% where AI is integrated more deeply across the lifecycle. In one client engagement, code migration work ran over 60% faster than baseline.



Engineering throughput:

In that same engagement, PRODIGY-enabled engineers produced 4.3 pull requests per week versus 2.9 for the baseline team (a 50% increase in throughput).



Operational cost:

PRODIGY's horizontal accelerators have reduced IT operations costs by 25% in live engagements, with no added headcount.



Time to production:

In a recent healthcare engagement, PRODIGY cut overall PDLC time by 40%, taking a platform from concept to processing 35 million data points daily for 11 million customers in 18 months.

How Organisations Can Start

Organisations do not need to adopt the full PRODIGY ecosystem at once.

The model allows enterprise teams to begin with the point of highest value or greatest urgency, then scale what works.

01 **Start with strategy**

For organisations that need to align leadership, define AI priorities, assess readiness, and identify the use cases most likely to create measurable value.

03 **Start with a focused accelerator**

For teams that already have a clear business problem and want to move faster by building from reusable AI capability rather than a blank slate.

02 **Start with a product opportunity**

For organisations that want to validate a new AI-enabled product, prototype quickly, and test whether the opportunity is worth scaling.

04 **Start with platform implementation**

For organisations ready to embed AI capabilities more broadly across products, workflows, systems, and delivery teams.

Summary

Ciklum PRODIGY provides the architecture, accelerators, and delivery framework required to build and operate AI-enabled systems at enterprise scale.

By combining reusable technology assets with structured delivery methods, PRODIGY helps organisations move from isolated AI experiments to scalable, governed, and business-ready AI systems.

PRODIGY helps organisations build AI capability that is not only faster to launch but also easier to govern, integrate, and scale across the business.

Explore how PRODIGY can support your AI roadmap

PRODIGY Team



James Lennon

Chief Technology Officer

✉ james.lennon@ciklum.com

“The biggest AI gains don't come from coding faster - they come from removing friction across the entire product development lifecycle. If AI is only improving your coding, you're optimizing a step, not your business. That's what PRODIGY was built for.”



Yannique Hecht

Senior Director Global AI R&D

✉ yannique.hecht@ciklum.com

“The difference between AI experimentation and AI transformation isn't the model - it's optimizing the entire product development lifecycle. That's the thinking behind PRODIGY, enabling organizations to become truly AI-ready with faster delivery, lower costs, and better engineering outcomes.”