What Nearly 70 GenAI Projects Taught Us About What Actually Works

Guide for business leaders
March 2026

Zürich
Switzerland

Guide for business leaders investing in generative AI — based on an in-depth analysis of real-world projects across Switzerland and Europe, many with global reach, each reviewed by senior industry experts in AI, cloud infrastructure, venture capital, and enterprise technology

What We Learned from 70 GenAI Projects

After studying nearly 70 generative AI projects — from five-person startups to organizations with tens of thousands of employees, across healthcare, insurance, agriculture, manufacturing, education, and public services — one pattern stood out above all others: the projects that created the most value went deep. Deep in one domain. Deep in one workflow. Deep in one measurable outcome. And they built the evidence to prove it.

What the Strongest Projects Got Right

The most intensive AI users are already six times more productive, and the gap is widening

Consider an insurance group that needed to translate regulated content across multiple languages. Rather than buying a generic tool, they locked legal terminology into the workflow and built human expert review into every step. The result: thousands of production requests, near-perfect reliability, and 80% less dependency on external translation services. Or consider a global consumer goods company that deployed an AI assistant across more than a hundred markets — and tracked adoption so rigorously they could report over 90% uptake and measurable hours saved per employee per week.

Now compare that with the many projects — some technically impressive — that couldn't answer a basic question: "What changed for the business?" Reviewers flagged this in more than three quarters of all projects. The technology worked. The measurement didn't exist.

That gap matters more now than it did a year ago. Industry data shows a widening divide between organizations that adopt AI deeply and those that experiment casually. The most intensive users are already six times more productive than the median — and the gap grows wider with more advanced capabilities[1]. Organizations that go deep early are building compounding advantages that late movers will struggle to close.

[1] OpenAI, *The State of Enterprise AI*, 2025 Report

01 Start with one workflow, go deep, then expand

Choose the single workflow with the highest volume of repetitive, knowledge-intensive tasks.

Define success criteria before selecting technology. Prove value in one place first, then broaden deliberately once you have evidence.

02 Define four metrics before building anything

Adoption, efficiency, quality, and cost.

Build instrumentation into the product so these are captured automatically from day one.

03 Design human oversight as a feature

Place human review at critical decision points and instrument the review process to capture quality data.

This generates training data for the next version and gives your compliance team the evidence they need to approve broader rollout.

04 Start governance documentation on day one

A data flow diagram, a retention policy, and an audit trail. These take a day to create and save months of stakeholder negotiations.

05 Invest in your proprietary knowledge advantage

Identify what you know about your domain that a general-purpose AI doesn't.

Encode that knowledge systematically into your AI systems. This is your moat.

06 Budget for adoption, not just deployment

Treat internal enablement as a product.

Identify AI champions per team, design tools to fit existing workflows, and track adoption by role.

If people aren't using it, fix the product, not the people

The most impactful projects didn't try to "bring AI to the enterprise." They picked a single, painful workflow and rebuilt it around AI, then measured what changed.

Suggestion 01

Start with one workflow and owned it completely

Real-World Example

A digital marketing company deployed a GenAI sales assistant for their SME clients. Rather than building a broad AI platform, they focused on one thing: helping sales teams write better out-reach emails.

They deployed in three months and measured a 25% increase in sales alongside 3x cheaper operations. The narrow focus made measurement straightforward and adoption natural — the tool did one thing, and it did it well.

Similarly, a healthcare technology company built AI specifically for clinical documentation in Swiss medical practices. They focused on the administrative burden doctors hate most — writing up consultations, generating billing codes, handling multilingual patient records. Within 16 months: 45 paying healthcare customers, 300+ active users, and a named hospital deployment. The key decision was resisting the temptation to build a general purpose healthcare AI and instead going deep on the documentation workflow that consumes hours of every physician’s day.

Across the dataset, this pattern held. Projects that focused on a single workflow consistently showed stronger adoption and clearer outcomes than those attempting broad transformation. Expert reviewers repeatedly praised "clear, focused use cases" and flagged "scope too broad" as a risk. But here's the nuance: once these teams proved value in their first workflow, the ones that expanded deliberately across task types saw their returns compound. Productivity gains don't just add up — they multiply[1].

What this means for you

Before building anything, identify the one workflow where your team spends the most time on repetitive, high-volume tasks that require domain knowledge. Start there. Measure time-per-task before and after.

Resist the urge to expand until you can prove value in that first workflow — but once you do, expand deliberately across task types. The evidence shows that productivity gains compound with breadth of use: teams that apply AI across multiple workflows report dramatically higher returns than those that stay in a single lane.

[1] OpenAI, *The State of Enterprise AI*, 2025 Report

The strongest projects didn't treat human oversight as a regulatory checkbox. They designed it as a core part of the product — and it made their solutions better.

Suggestion 02

Keep humans in the loop and made it a feature, not a compromise

Real-World Example

A pharmaceutical company built an AI-powered patient support system where a clinical safety layer monitors every interaction in real time, ready to intervene.

The result wasn't just compliance — it was complete patient satisfaction in their clinical study, with healthcare professionals reporting that the AI enhanced rather than replaced their care.

An insurance company processing legal documents took a different but equally deliberate approach. They required human review of every AI-generated case summary and built traceable outputs, evaluation sets, and data drift monitoring into the system from day one. What could have been a bottleneck became a selling point internally, accelerating adoption rather than slowing it.

This wasn't coincidental. Expert reviewers consistently praised projects that treated governance as an enabler. The projects with the strongest trust postures — Swiss data sovereignty, audit trails, human approval gates — were also the ones with the highest adoption rates. Governance didn't slow them down; it gave stakeholders the confidence to say yes.

What this means for you

Design your AI workflow so that humans review outputs at critical decision points — and instrument that review process to capture quality data. This isn't overhead.

It's your training data for the next version, and it's the evidence your compliance team needs to approve broader rollout.

The projects with the strongest business cases didn't add measurement as an afterthought. They designed their AI systems to generate evidence of their own value from the first deployment.

Suggestion 03

Measure from day one and made measurement the product

Real-World Example

A global enterprise that deployed AI across more than a hundred markets and tens of thousands of employees set itself apart not with the technology but with the measurement infrastructure: adoption rates by market, hours saved per employee per week, compliance incidents, and reuse rates across components. When leadership asked "is this working?", they had production data — not estimates, not projections, but answers.

A digital commerce company took the same approach at a smaller scale, instrumenting their AI assistant across hundreds of customer service agents and measuring a multi-hour reduction in ticket resolution time. The numbers were modest but real — and because they were measured rigorously, they justified expansion to additional departments.

Contrast this with the many projects that reported impressive capabilities but couldn't quantify their impact. Expert reviewers flagged this pattern repeatedly: "No quantified business impact." "Claims presented without methodology." "Impact metrics are projections, not measurements." Even technically excellent projects struggled to secure continued investment when they couldn't demonstrate value in business terms.

What this means for you

Define four metrics before you write a single line of code:

- adoption (are people using it?)
- efficiency (is it saving time?)
- quality (is the output good enough?)
- cost (what does it cost per unit of work, and how does that compare to the baseline?)

Instrument your system to capture these automatically. The measurement infrastructure is as important as the AI itself.

The projects that created lasting competitive advantage didn't just deploy off-the-shelf AI. They combined general-purpose models with proprietary domain knowledge that competitors couldn't easily replicate.

Suggestion 04

Build on proprietary knowledge, not generic capabilities

Real-World Example

An agricultural technology company built a marketing platform with specialized AI agents trained on proprietary data from global operations.

The agents understood crop cycles, regional regulations, and farmer communication preferences — knowledge a generic AI tool couldn't match.

What this means for you

That knowledge — your data, your terminology, your workflows, your edge cases, your regulatory requirements — is your moat.

The team that spends three months encoding domain expertise into their AI will outperform the team that spends three months evaluating which foundation model to use.

What do we know about our domain that a general-purpose AI doesn't?

Where Most Projects Had Room to Grow

The evidence gap

The most common challenge wasn't technical — it was evidentiary. Project after project built impressive systems but couldn't demonstrate their business value because they hadn't set up measurement from the start.

One project processed tens of thousands of financial documents daily at genuine enterprise scale. Yet expert reviewers asked for clearer articulation of what the AI actually contributed versus traditional automation. The technology worked. The evidence infrastructure didn't exist.

In a world where every vendor claims AI-powered everything, the ability to prove your results is the differentiator. Not "we think it's faster" — but "we measured a multi-hour reduction in resolution time across hundreds of agents over two months." Treat evidence infrastructure as a first-class requirement. Before your first pilot, define what success looks like in numbers. During the pilot, capture baselines. After deployment, report results with methodology, timeframe, and sample size. A single rigorous case study — with real numbers on adoption, efficiency, quality, and cost — is worth more than a hundred capability demonstrations.

The ability to prove your results is the differentiator

The governance gap

The second most common gap was the absence of documented ethics, privacy, and governance frameworks — even in projects handling sensitive data in regulated industries.

Projects that handled patient data, financial documents, or personal information often had no documented approach to data privacy, no explicit consent frameworks, and no audit trails. This wasn't because the teams didn't care — it was because governance was treated as something to add later, after the technology was proven.

The counter-examples told a different story. An insurance company built their AI case management system with DSG[2] compliance, traceable outputs, and data drift monitoring from the start. A healthcare screening startup obtained CE Class I medical device marking and built Swiss sovereign infrastructure before scaling. These projects didn't move slower — they moved with more confidence, because every stakeholder conversation started from a position of trust rather than risk.

Start your governance documentation on day one, not after launch. You need three things: a data flow diagram showing where information goes, a retention policy explaining how long you keep it, and an audit trail showing who accessed what. These documents take a day to create and save months of stakeholder negotiations later.

Start your governance documentation on day one, not after launch

[2] DSG: Datenschutzgesetz, the Swiss Federal Act on Data Protection

The focus gap

A recurring pattern was projects with ambitious, multi-capability visions that struggled to demonstrate depth in any single area.

One project combined a knowledge graph, personalized AI enrichment, a digital twin, a podcast studio, and a data marketplace — all in a single platform.

The expert feedback was direct: "Pick the single feature that generated the strongest reaction during demos and build a focused go-to-market around it." The projects that avoided this trap shared a common discipline: they said no to good ideas in order to execute great ones.

If your AI initiative has more than three capabilities on the roadmap for the first six months, you're probably trying to do too much. Pick the one that solves the most painful problem — whether that's an internal bottleneck, a customer-facing friction, or an untapped opportunity. Prove it works. Then expand.

Pick the single feature that generated the strongest reaction and build a focused go-to-market around it.

Adoption followed from design, not from training decks.

The skills gap

A pattern that cut across nearly every project — but was rarely addressed explicitly — was the gap between what the technology could do and what the organization was ready to absorb. Projects launched with strong technical foundations but without structured plans for training, change management, or internal capability building.

As organizations move toward more autonomous AI systems, the ability to supervise, collaborate with, and strategically direct AI — what some call "agentic literacy" — is becoming as important as traditional digital literacy[3]. Even among active daily AI users in large enterprises, significant percentages have never tried the most capable tools available to them — not because the tools don't exist, but because nobody showed them how or why[1].

The strongest projects addressed this by design. The global consumer goods company that achieved over 90% adoption treated internal enablement as a product in its own right. The insurance company that built AI into their core claims platform designed the integration so that using the AI was the default workflow, not an optional add-on.

Budget for change management as a line item, not an afterthought. Identify AI champions in each team who can model usage and coach peers. Design your AI tools to fit into existing workflows rather than requiring users to adopt new ones. And track adoption by team and role — if a department isn't using the tool, that's a product problem, not a people problem.

[1] OpenAI, *The State of Enterprise AI*, 2025 Report

[3] AWS, *Agentic AI for Leaders: A Practical Guide for Executives*, 2025

What to Do Next

What separates the strongest GenAI projects

If you haven't started yet

- Identify your highest-pain workflow.
- Interview five people who do this work daily.
- Measure how long it takes today.
- Define your four metrics.

If you're piloting

- Run a 30-day measurement cycle with baselines captured on day one.
- Report results with methodology and sample size on day 30.
- Build one rigorous case study with real before-and-after numbers

If you're ready to scale

- Document your governance framework and get compliance sign-off.
- Invest in change management and internal enablement.
- Expand to adjacent workflows based on evidence, not ambition.

The technology is ready

The question is whether your organization can prove it works, govern it responsibly, and focus long enough to let it deliver real value. The projects in this analysis show that the ones who did all three created outcomes that speak for themselves.

About This Guide

This guide was produced as part of the GenAI Zürich 2026 awards program. The analysis is based on nearly 70 project submissions across three award tracks — Rising Innovators, Impact Achievers, and Enterprise Transformers — each independently scored and reviewed by a panel of senior industry experts. All company names and identifying details have been anonymized.

How this guide was created

The analysis and initial drafting of this guide were produced using an agentic AI workflow, with expert jury members reviewing and commenting on each version. Their feedback was fed back into the process, making this a human-in-the-loop collaboration between AI and domain experts.

Kiro CLI, an agentic AI development tool by AWS, served as the orchestration environment. Custom Agent SOPs (Standard Operating Procedures from the Strands Agents open-source project) defined the end-to-end workflow — from ingesting all submissions and expert scores, to cross-referencing judge feedback, identifying patterns across tracks, and generating the final text.

The underlying model was Anthropic Claude (Opus 4.6).
For an introduction to this agentic tooling, visit goagentic.ch

Data Sources

Industry data referenced in this guide draws on published research from
OpenAI "The State of Enterprise AI", 2025 Report,
Anthropic "Building Trusted AI in the Enterprise", and
AWS "Agentic AI for Leaders: A Practical Guide for Executives"

Expert jury

Marc Stampfli
Business Director Switzerland
@ NVIDIA

Stefan Ravizza
Co-Founder
@ Artifact

Dr. Efi Pylarinou
Global Fintech & Tech Thought Leader
@ GrowFin

Claudia Pletscher
Entrepreneur, Board Member & Innosuisse Innovation
@ Council

Sacha Ghiglione
AI & Strategy Executive

Paul Meyrat
AI Officer
@ Stadt Zürich

Anne-Marie Buzatu
Executive Director
@ ICT4Peace Foundation

Dr. sc. ETH Michael Hardegger
Lead Data Scientist
@ Digitec Galaxus AG

Dr. Sina Wulfmeyer
Chief Data Officer
@ Unique AI

Dr. Jan Kerschgens, EMBA
Innovation & DeepTech Leader in AI

Daniel Dobos
Research Director
@ Swisscom

Alex Dean
Manufacturing Practice Lead
@ Unit8

Rea Selina Wenk
CDO
@ CSS

Dr. Denis Samuylov
Co-Founder
@ GenAI Zürich

Expert jury and guide production

Christoph Schnidrig
Head of Technology
@Amazon Web Services
Orchestrated and iterated the agentic AI workflow that produced this guide

Michael Wegmüller
Co-Founder
@Artifact
Acted as lecturer and editorial advisor