Managing risks and cultural change to overcome AI

Version 04.03.2026

Content

Possible Strategies 2
A) Waiting for market readiness 2
B) Quickly implement selective pilots (quick wins) with or without partners 2
C) Strategic competence building for sustainable advantages (focused solutions) .. 3
Four critical areas of action 3
1. Governance & Compliance (The Shield) 4
2. Vibe Coding & Tech Stack (The Workbench) 4
3. Core AI Tech & Business Use (The Engine) 5
4. Internal Utilities & Standards (The Operating System) 5
Conclusion & Call-to-Action: Transitioning to an Active AI-Driven Organization 6
1. Structural Clarity: Governance as a Competitive Advantage 6
2. The Strategic Pillar: Data Protection & Contracts 7
3. Cultural Transformation: From Passive to Active Leadership 7
4. Management Call-to-Action 7
5. Measuring Success: The KPI Framework for AI Transformation 8
Proposed Management Mandate: The Next Steps 8
Appendix: Deep Dive Governance, Compliance & Technology Synergy 9
Appendix: Deep Dive – Vibe Coding & Tech Stack 2026 11
Appendix: Deep Dive Core AI Technology & Business Landscape 2026 13
Appendix: Deep Dive Shadow AI 17

Possible Strategies

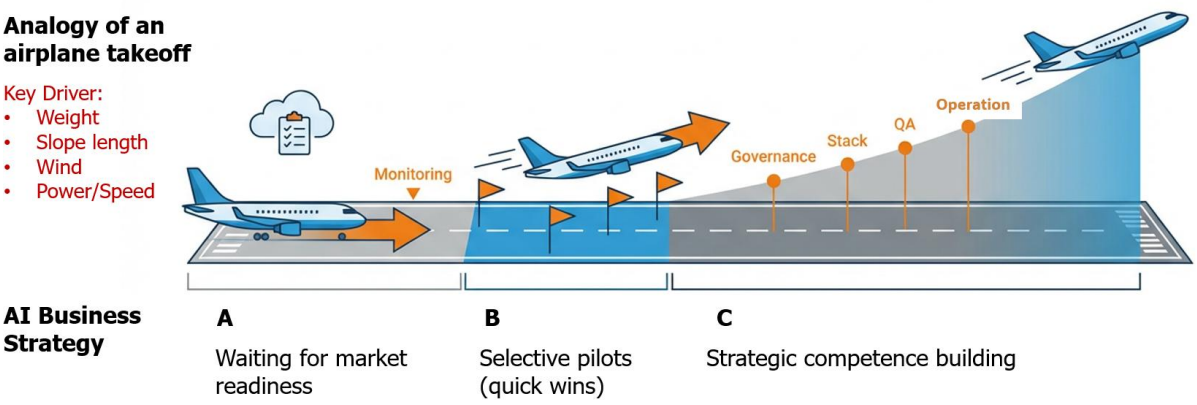

A) Waiting for market readiness

- **Core idea:** We observe the market and only invest substantially once technologies, platforms, providers, and best practices are sufficiently stable.
- **When it makes sense:** When compliance/liability risks are high, requirements are still unclear, or the business case is not yet robust.
- **Advantages:** Low initial investment, less risk of making the wrong decision, use of proven standards and cheaper components later on.
- **Disadvantages/risks:** Missed learning curve, potential reputation and time-to-market disadvantage, dependence on external roadmaps.
- **Consequences (ownership/organization):** Minimal competence building, focus on monitoring, defining criteria for when "market readiness" is achieved.

B) Quickly implement selective pilots (quick wins) with or without partners

- **Core idea:** We quickly deliver 1–3 concrete pilots to achieve measurable market or efficiency advantages early on and build real learning curves.
- **How it works:** Selection of clearly defined use cases with short lead times, defined success measurement, and controlled risk zone (e.g., UAT or selective production).
- **Role of partners:** Partners can provide speed and specialist expertise, but pdc retains clear product and data sovereignty (and defines exit options).
- **Advantages:** Rapid validation, internal momentum, early visible results, learning curve for processes/compliance/QA.
- **Disadvantages/risks:** Risk of "pilot fragmentation," technical debt, dependence on partners, lack of reusability without standards.

- **Consequences (ownership/organization):** Lightweight delivery setup, clear pilot governance, defined criteria for scaling or stopping.

C) Strategic competence building for sustainable advantages (focused solutions)

- **Core idea:** We build targeted internal capabilities (governance, architecture, delivery, QA, operations) to integrate AI solutions into solutions in a scalable, compliant, and repeatable manner.
- **When it makes sense:** When AI is to become a central component of value creation, differentiation through quality/integration arises, and long-term ownership is required.
- **Advantages:** Independence, consistent quality, better margin prospects, reusability (standards/assets), clear responsibilities.
- **Disadvantages/risks:** Higher initial investment, longer time-to-value, recruiting/enablement effort, need for consistent governance.
- **Consequences (ownership/organization):** Establishment of a clear operating model structure (roles, decision rights, standards) and a focused backlog with prioritized solutions.

Four critical areas of action

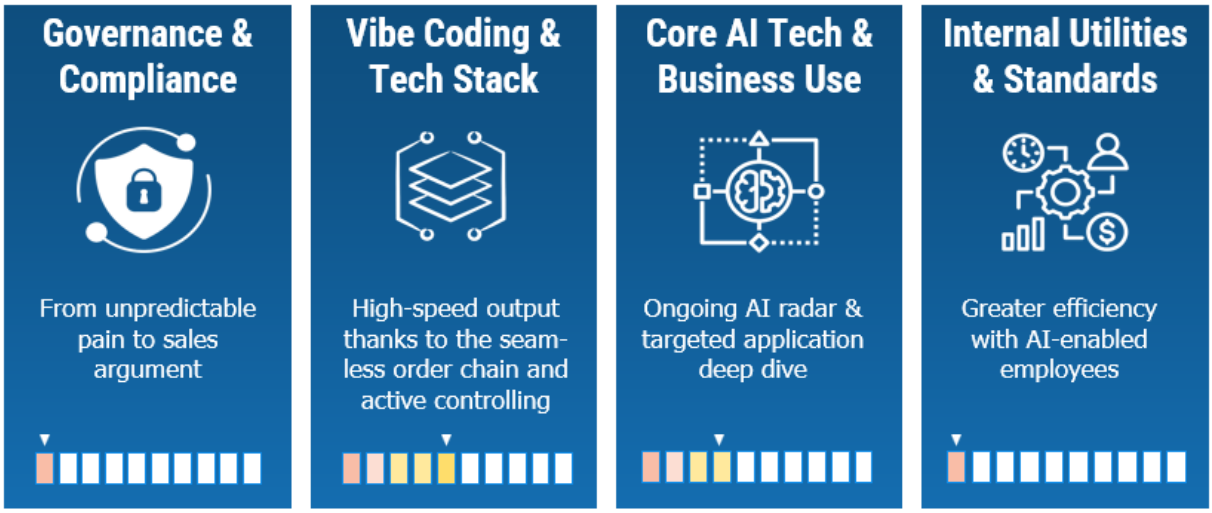

1. Governance & Compliance (The Shield)

See background infos at **Appendix: Deep Dive Governance, Compliance & Technology Synergyworst-case scenarios**

worst-case scenarios

- **Exclusion from major tenders:** Without proactive EU AI Act certification, pdc will lose its qualification for pitches at banks, insurance companies, or IT corporations that require seamless AI governance from their service providers starting in 2026.
- **Data exfiltration through "shadow AI":** If we do not set enterprise standards, employees will continue to upload sensitive CRM data or confidential strategy papers to free tools, resulting in pdc IP being used to train public models and leaving internal control forever.

Typical questions in the real world today

- Have we listed the tools used in the data protection documents of the customer contracts and documented other legal safeguards?
- Do we comply with the requirements of the EU AI Act?

2. Vibe Coding & Tech Stack (The Workbench)

See background infos at **Appendix: Deep Dive – Vibe Coding & Tech Stack 2026**

worst-case scenarios

- **Competition dependency due to inertia:** While competitors can "vibon" individual lead workflows in 48 hours using Vibe Coding, pdc needs weeks to make the same adjustments in the classic development loop – we are being overtaken in terms of price and time.
- **Uncontrollable technical debt:** Without a complete documentation chain (client > implementer > user), "one-person AI solutions" are created that no one can maintain once the person who created them leaves the company.

Typical questions in the real world today

- We use Cursor AI for vibe coding. Is this the best option, or would First Mover and others like Claude Code be better? Who is keeping an eye on the constantly changing landscape?
- Have we set up a "security-first" configuration in Cursor? Have we integrated tools that analyze the generated code immediately? Have we established review processes and testing procedures as standard? All this to ensure that the resulting code is checked in terms of quality and security?

- Have we set up the optimal documentation workflow PRD > Vibe Coding > Technical and training documentation?

3. Core AI Tech & Business Use (The Engine)

See background infos at **Appendix: Deep Dive Core AI Technology & Business Landscape 2026**

worst-case scenarios

- **Strategic blindness:** Without continuous radar, we will notice too late when leading-edge AI replaces our core product (such as parts of xRelation) with a simple, cheap API solution – we will then be investing in dead technology.
- **Efficiency gap to the market:** If we stick with prototypes instead of deeply integrating agentic workflows, pdc will remain a "manufacturer," while competitors operate their lead management as a highly scaled, AI-driven "factory" at a fraction of the cost.

Typical questions in the real world today

- Should we commit to OpenAI at all, given that Claude is currently leading in coding and Gemini in skill services?
- Do our applications have to be run with LLM, or could we run them more cost-effectively or even locally with SLM and train them specifically for our needs?

4. Internal Utilities & Standards (The Operating System)

See background infos at

Appendix: Deep Dive Shadow AI

worst-case scenarios

- **Talent drain & loss of motivation:** High-potential employees at pdc will leave the company if they work with AI power in their private lives but are held back in their productivity at the office due to a lack of subscriptions and rigid prohibitions.
- **Cultural stagnation:** If the broad workforce is not empowered now, a two-tier society will emerge between those who are “AI-savvy” and the rest – this will paralyze the entire organization and lead to massive internal resistance to future changes.

Typical questions in the real world today

- Should we use chatGPT subscriptions or build an internal GUI for the existing connection via MS Azure (secure, hosted in Zurich and Frankfurt) to OpenAI GPT-5.2? Can such a GUI also offer the functions of projects or GPTs as well as user-specific personalization?
- Which AI tools are permitted for employees to use, and which subscriptions are paid for by the company? Which employees are allowed to use AI tools on an experimental basis?
- Which utility tools may employees use? For example, speech-to-text, image generation, or AI-supported productivity tools such as Notion AI or NotbookLM, etc.?

Conclusion & Call-to-Action: Transitioning to an Active AI-Driven Organization

The implementation of the **AI Measures and Organizational Structure** presented here is not merely a technical upgrade; it is a fundamental shift in our corporate operating system. We are moving beyond the stage of isolated experimentation toward a **scalable, industrial engine for innovation**.

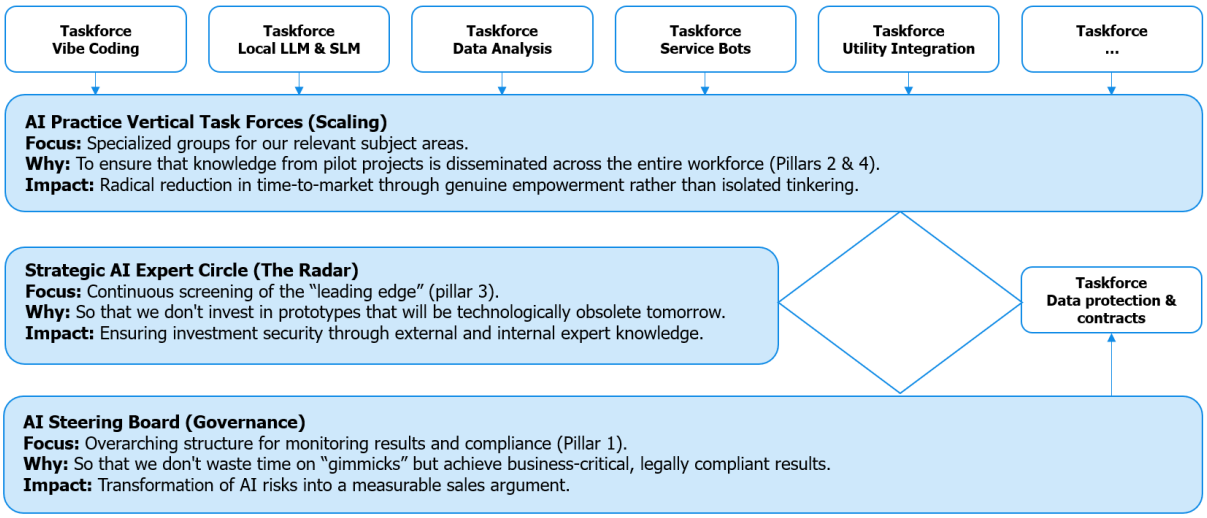

1. Structural Clarity: Governance as a Competitive Advantage

To combine speed with enterprise-grade security, we are establishing a clear hierarchy of accountability:

- **AI Steering Board (The Governance):** Provides the overarching framework to monitor results and ensure compliance. This ensures we don't lose time on "gadgets" but deliver business-critical, legally sound results.
- **Strategic AI Expert Circle (The Radar):** Protects our investments. By continuously screening the "Leading Edge," we ensure that we do not invest in prototypes that will be technologically obsolete tomorrow.
- **Vertical Taskforces (The Scaling):** This is where specialized groups—from *Vibe Coding* to *Service Bots*—ensure that knowledge from pilot projects is integrated into the broader workforce, drastically reducing time-to-market.

2. The Strategic Pillar: Data Protection & Contracts

A critical component of this structure is the **Taskforce for Data Protection & Contracts**. In the AI era, compliance is no longer a "bottleneck" but a **measurable sales argument** and a prerequisite for trust.

- **Risk Transformation:** This taskforce is responsible for transforming AI-related risks into managed assets.
- **Legal Readiness:** By proactively managing contracts and data privacy standards, we create the "safe harbor" that allows our operational taskforces to innovate at high speed without legal hesitation.

3. Cultural Transformation: From Passive to Active Leadership

Structure alone is insufficient. We require a cultural shift that must be understood and championed by management (Top-Down) while being anchored deeply within the teams (Bottom-Up).

- **Active Ownership:** We must move from a passive, observant stance to a highly proactive one. Leadership must not only "permit" AI but actively drive the adoption of AI-tailored workflows.
- **Shifting the Focus of Work:** Our focus is shifting from "Execution" to "Orchestration." As AI agents increasingly take over the creation of objects and elements, management and staff must invest more time in **setting precise objectives, defining standards, and performing rigorous quality control.**
- **Data Discipline:** A structured approach to data is now a mandatory standard. We must lead by example in establishing these new ways of working.

4. Management Call-to-Action

We hereby call for the formalization of these Taskforces under the guidance of the Expert Circle and the Steering Board.

Your Mandate:

As managers, your role is to ensure this structure is lived. Bridge the gap between technical potential and business value. Empower your teams to use these Taskforces as catalysts for growth.

The goal is clear: **Less manual execution, more strategic steering.** By embracing this structured AI framework, we secure our investment safety and transform AI risks into a decisive market advantage.

5. Measuring Success: The KPI Framework for AI Transformation

To ensure that this structure delivers on its promise of "measurable sales arguments" and "investment safety," we will track progress across four key dimensions:

Dimension	**Key Performance Indicator (KPI)**	**Target / Goal**
Operational Efficiency	**Task Automation Ratio**	Reduction of manual execution time in core tasks by **25%** within the first 12 months.

Dimension	Key Performance Indicator (KPI)	Target / Goal
Speed-to-Market	**Lead Time from Pilot to Scale**	Reducing the time it takes to move a validated AI prototype through the *Vertical Taskforces* into production by **40%**.
Governance & Trust	**Compliance Readiness Score**	100% of AI projects must pass the *Taskforce Data Protection & Contracts* audit before scaling.
Strategic Edge	**"Leading Edge" Adoption Rate**	Percentage of active projects utilizing technologies identified as "top-tier" by the *Strategic AI Expert Circle*.
Cultural Shift	**AI Literacy & Self-Service Rate**	Percentage of staff actively using internal AI tools/LLMs for daily tasks without external IT support.

Proposed Management Mandate: The Next Steps

As a leadership team, your immediate focus should be on the following:

1. **Resource Allocation:** Formally appoint the leads for each **Vertical Taskforce** and ensure they have the "airtime" to implement standards.
2. **Define the "North Star":** Work with the **AI Steering Board** to define the top three business-critical goals for the next quarter.
3. **Lead by Example:** Shift your own reporting and decision-making processes to focus on **Objectives & Controls** rather than micro-managing the creation of deliverables.

Final Note: The *Taskforce Data Protection & Contracts* is your safety net. By empowering them to set standards early, you are not slowing down—you are building the high-speed tracks that allow the rest of the organization to run at full velocity.

Appendix: Deep Dive Governance, Compliance & Technology Synergy

1. Initialization of Oversight Structures

Before the operational taskforces begin their work, the following two bodies must be formally established as a regulatory foundation:

- **AI Steering Board (Strategic Instance):**
  - **Composition:** Executive Board, Head of IT, Data Protection Officer (DPO).
  - **Core Task:** Definition of the "Red/Green List" for approved AI tools and final approval of high-risk applications according to the EU AI Act.
  - **Management Focus:** Transforming compliance risks into a competitive sales argument ("Trusted AI by pdc").
- **Strategic AI Expert Circle (Radar & Audit):**
  - **Composition:** Internal AI champions and external subject matter experts.
  - **Core Task:** Technical screening of the "Leading Edge". Ensuring that chosen technologies (e.g., Agentic Workflows) remain technically auditable and do not represent "black boxes" without access control.

2. Technology & Data Protection: An Inseparable Unit

In our previous "GEMS AI-Strategy" discussions, we established that technology decisions directly determine legal maneuverability. In 2026, data protection is no longer an "add-on" but an integral part of the tech stack.

A. US Cloud vs. Swiss Sovereignty (The Location Effect)

Technologies like **OpenAI or Gemini** offer maximum performance but are subject to the **US CLOUD Act**.

- **Implication:** For international corporations with Swiss subsidiaries, the **GDPR** often applies (Principle of Extraterritoriality).
- **Strategic Solution:** Sensitive customer data is processed via **Local LLMs/SLMs** (such as Gemma 2 or Mistral) running on our own infrastructure. This ensures that data sovereignty physically remains in Switzerland and fully complies with the Swiss **FADP (DSG)**.

B. Vibe Coding & Code Sovereignty

The use of **Natural Language Programming (Cursor/Windsurf)** massively accelerates development.

- **Privacy Aspect:** AI-generated code must undergo automated **security scans**. We thus prevent the AI from unintentionally introducing security gaps or "hardcoded credentials" that would violate the accountability obligations of the GDPR.

C. Agentic Workflows & Transparency Obligations (EU AI Act)

The deployment of specialized AI agents for lead scoring falls under the **transparency and documentation obligations** of the EU AI Act.

- **Technology Nexus:** We are implementing an **AI API Gateway** that logs every interaction. This fulfills the legal requirement for **auditability** and protects pdc from liability risks arising from AI-driven decisions.

3. Compliance as a Competitive Advantage (USP)

For international corporations, pdc is not just a technical implementer but a **"Compliance Enabler"**.

- Through a **seamless chain of custody** (from prompt to output), we ensure that no uncontrolled leakage of proprietary IP occurs into public models.
- Proactive alignment with the **EU AI Act** secures market access across the EU, while the Swiss **FADP (DSG)** forms the basis for protecting the personal rights of our local clients.

Appendix: Deep Dive – Vibe Coding & Tech Stack 2026

1. Executive Summary: What is Vibe Coding?

In traditional programming, a developer writes specific lines of code in languages like Python or JavaScript to tell a computer exactly what to do. **Vibe Coding** (Natural Language Programming) flips this script. The developer (or project manager) describes the "vibe"—the intent, logic, and desired outcome—in plain English or German.

At pdc, we use **Cursor AI** as our standard workspace (IDE). Cursor doesn't just suggest words; it understands our entire codebase and generates functional software components based on our descriptions. This allows us to move from an idea to a working prototype up to **10 times faster** than before.

2. The Tech Stack: Our Digital Workbench

To maintain leadership in lead management, our "workbench" is built on three layers:

- **The Interface (Cursor AI):** Our primary environment where human intent meets machine execution.
- **The Engine (KI-API-Gateway):** A centralized "control center" that connects our tools to the world's most powerful AI models while keeping our data private.
- **The Product (xRelation & Agentic Workflows):** Our CRM and lead-scoring tools that are now being enhanced by specialized AI agents rather than just static code.

---

3. Risk Analysis: The "Vibe Code Debts"

While speed is a competitive advantage, "vibe-driven" development introduces specific risks that we categorize as **Vibe Code Debts**.

Risk Factor	Definition	Business Impact
Technical Debt	AI-generated code that works "now" but is messy or hard to fix later.	High maintenance costs and system crashes in 2027/2028.
Shadow AI	Employees using unapproved, "free" AI tools for work.	Data leaks; proprietary customer data ending up in public models.
Code Wildgrowth	Too many apps/features created too fast without a central plan.	A "spaghetti" system that no one understands or can audit.
Compliance Gaps	Code that violates the **EU AI Act** or data privacy (DPIA).	Legal sanctions, fines, and loss of "Trusted Partner" status.

---

4. Strategic Guardrails: How We Scale Safely

To prevent these risks from becoming "pains," pdc has implemented a "Shield and Motor" strategy:

A. The AI Steering Board (Governance)

We do not code in a vacuum. Every major AI implementation is overseen by the **Steering Board** to ensure it aligns with business goals and legal requirements. We transform compliance from a "boring rule" into a **Sales USP**—we can prove to enterprise clients that our AI is safer than the competition's.

B. Security-by-Design & Code Audits

We don't just "trust the vibe." We implement **Code Audits**. This means AI-generated code is reviewed by human experts or specialized "Guardrail AI" to ensure it follows our security standards before it goes live.

C. Lean Documentation Workflows

Since AI writes the code, we use AI to write the **documentation** simultaneously. This ensures that even if a developer leaves, we know exactly how the "vibe" was turned into a "feature".

---

5. The Taskforce Vibe Coding: Your Rapid Response Team

Within our new **AI Practice Vertical**, the **Taskforce Vibe Coding** has one mission: **Enablement without Chaos**.

- **Mission:** Move knowledge from "isolated tinkering" into company-wide standards.
- **Impact:** They ensure that Marketing and Project Management can adjust the CRM independently of IT bottlenecks, but always within the "safe harbor" of our Tech Stack.

---

6. Conclusion for Management

Vibe Coding is our **High-Speed-Output** engine. By using **Cursor AI** and our **API-Gateway**, we eliminate technical bottlenecks. However, our leadership depends on **Governance**. We are not just "playing" with AI; we are building a professional, certifiable infrastructure that scales without increasing headcount.

Key Takeaway: We accept the "Vibe" for speed, but we enforce the "Standard" for safety.

Appendix: Deep Dive Core AI Technology & Business Landscape 2026

Introduction: The Shift from Hype to Infrastructure

As of 2026, Artificial Intelligence has moved beyond a "productivity tool" to become the core infrastructure of modern business. For the management board, the challenge is no longer "if" we use AI, but how we architect our AI stack to ensure **sovereignty, cost-efficiency, and competitive agility**. This appendix provides a comprehensive explanation of the technical landscape and its business implications.

---

1. The Geopolitical Data Dilemma (US Dominance)

Most leading AI models (Large Language Models or LLMs) like ChatGPT (OpenAI), Gemini (Google), and Claude (Anthropic) are developed in the United States.

- **Simplified Explanation:** Think of an LLM as a "digital brain." Currently, the most powerful brains are American. However, the US legal framework for data privacy is significantly more relaxed than the Swiss or European (GDPR) standards.
- **The Risk:** When your employees or customers interact with these models, data is processed in a legal environment where privacy is not the primary priority. Even if the provider promises not to "train" on your data, the metadata and interaction logs remain within US jurisdiction.

2. Sovereignty & The "Cloud Act" (The Myth of Local Storage)

Many US providers offer hosting on servers physically located in Switzerland or Europe (e.g., Microsoft Azure's Swiss Data Centers).

- **The Reality Check:** The **US CLOUD Act** allows US federal law enforcement to compel US-based technology companies to provide requested data, regardless of where that data is stored globally.
- **Business Impact:** If our business handles highly sensitive client data (financial, medical, or proprietary R&D), hosting it with a US provider—even on Swiss soil—does not offer 100% protection against foreign government access. This is a critical factor for our "Swiss-Made" brand promise.

3. The Chinese Open-Source Paradox

China has become a powerhouse in AI, often releasing models that are "Open Source" (free to use and modify).

- **Why it Matters:** These models (like those from Alibaba or DeepSeek) are technically brilliant and often perform better than Western models in specific logic tasks.
- **The Risk:** While the code is open, the "weights" (the internal logic) are a black box. Using Chinese tech introduces a different layer of geopolitical risk and potential "backdoor" vulnerabilities that could compromise long-term security.

4. The "Feature Gap": Consumer vs. Enterprise AI

There is a massive difference between the ChatGPT app on your phone and the version we build into our business processes via Microsoft Azure or AWS.

- **The Problem:** The "Enterprise" versions often lag 3–6 months behind the consumer versions in terms of features.
- **The Effort:** Building an enterprise solution is not "plug-and-play." It requires a massive internal effort to build "skills" (orchestration) so the AI can actually talk to our CRM, our files, and our staff. It is a software engineering project, not just an "on" switch.

5. The Swiss Alternative: Sovereignty as a Service

Switzerland has recognized the risk of dependency and created its own AI ecosystem.

- **Key Players:** * **AlpineAI (SwissGPT):** A solution designed for Swiss companies that ensures data never leaves the country.
    - **Swisscom AI Platform:** High-performance hardware (NVIDIA SuperPODs) located in Switzerland.
    - **Apertus (ETH/EPFL):** A domestic, high-quality model trained with Swiss cultural and linguistic nuances (including Swiss German).
- **Strategy:** These are our "Safe Havens." They might be slightly more expensive, but they offer the legal and ethical "Swissness" that US providers cannot match.

6. Small Language Models (SLMs) & On-Device AI

We previously experimented with local LLMs and found the performance lacking. That has changed.

- **The New Generation:** Models like **Gemma 2** or **Mistral** are "Small Language Models." They are compact enough to run on a high-end laptop or even a smartphone without an internet connection.
- **FunctionGemma:** This is a breakthrough where a small model is specifically trained to "do things" (call functions, book meetings) rather than just "chat."

- **The Opportunity:** For many tasks, we don't need a "super-brain" in the cloud. A "mini-brain" running locally is faster, 100% private, and costs zero in monthly token fees.

7. The "Hidden Bill": Cost Complexity & Scaling

AI pricing is currently a "black box" that can lead to "sticker shock."

- **Fragmented Pricing:** You pay for "tokens" (roughly words), for the server being "awake," for data moving in and out, and for the engineers to keep it running.
- **The Volatility Risk:** If the "AI bubble" bursts or energy prices for data centers (which are enormous) double, the cost of running our AI Use Cases could skyrocket overnight.
- **Management Logic:** We must build our Use Cases to be "Model Agnostic"—meaning we can swap a expensive model for a cheaper one without rebuilding the whole system.

8. AI Orchestrators: The "Command Center" Strategy

Instead of picking one AI (like OpenAI), many companies use an **Orchestrator** (e.g., **Manus, LangDoc, or WAPI-AI**).

- **How it works:** Imagine a "universal remote control" that can talk to any TV. The Orchestrator allows us to use OpenAI for writing, ElevenLabs for voice, and a Swiss model for data privacy—all in one workflow.
- **The Pros/Cons:** It prevents "Vendor Lock-in" (being stuck with one provider). However, it adds a new layer of costs and a new "middleman" that we must trust.

9. Multimodality: The Case of VoiceBots

Our VoiceBot doesn't just use an LLM. It uses a chain of specialized tools:

1. **Ear:** Speech-to-Text (Transcribing the caller).
2. **Brain:** The LLM (Understanding the intent).
3. **Voice:** Text-to-Speech (e.g., **ElevenLabs** for a natural human sound).

- **The Challenge:** Each of these tools might run on a different server in a different country. Managing the "latency" (the delay in response) and the "data trail" across these different tools is the primary technical hurdle in 2026.

10. Niche vs. General Purpose (The "Sniper" Approach)

A common mistake is trying to use one "Big AI" for everything.

- **The Niche Strategy:** For specific departments (e.g., Legal or Accounting), it is often better to use a "Niche" solution—a small, highly specialized tool that does only one thing but does it perfectly.
- **The Advantage:** These tools are easier to audit, harder to "hallucinate" (lie), and significantly cheaper to run than a massive general-purpose model.

Executive Decision Guidelines for 2026

1. **Prioritize "Model Agnosticism":** Never build a solution that *only* works with OpenAI. Ensure we can switch providers within 48 hours if prices rise or compliance rules change.
2. **The 80/20 Privacy Rule:** 80% of tasks can use standard US cloud tools. The 20% involving core IP or customer secrets **must** stay on Swiss or Local (SLM) infrastructure.
3. **Monitor the "Total Cost of Ownership" (TCO):** Do not just look at the subscription price. Include the cost of the engineers required to maintain the "Orchestration" layer.
4. **Invest in "Human-in-the-Loop":** As we move toward "Agentic AI" (AI that acts on its own), our role shifts from "doing the work" to "supervising the AI's decisions."

Appendix: Deep Dive Shadow AI

When employees are faster than the IT department. The term shadow AI describes the use of AI tools by employees without the knowledge, approval, or control of the IT department.
by Roger Basler de Roca

TLDR: Over half of knowledge workers in the DACH region use AI tools that have not been approved by their company. In Switzerland, 97% of IT managers surveyed say they are confronted with shadow AI. The motives are understandable: time savings, productivity, lack of official alternatives. The risks, however, are considerable – from GDPR violations and data leaks to fines under the EU AI Act. Bans alone are not enough. What is needed are clear AI guidelines, secure internal platforms, and training. Shadow AI is not a marginal phenomenon – it is a seismograph for organizational action.

What is shadow AI and why might it affect your company?

The term shadow AI describes the use of AI tools by employees without the knowledge, approval, or control of the IT department. The term is based on the well-known shadow IT, but goes one step further: While shadow IT involves the use of an unauthorized project management tool, for example, shadow AI involves potentially sensitive company data flowing directly into external AI models. Think of employees who have contracts summarized in ChatGPT, copy customer data into a translation tool, or send application documents through an AI analysis – all without IT approval.

Shadow AI is already here: the figures speak for themselves

The data for 2024 and 2025 is clear. In its study “Chasing Shadows,” Software AG found that over half of all knowledge workers in the US, UK, and Germany use AI tools not provided by their companies. 75% already use AI, and the study predicts an increase to 90% in the near future.

The Bitkom 2025 survey (604 companies in Germany with 20 or more employees) confirms the trend: in 8% of companies, private AI use for professional purposes is “widespread” – double the figure for 2024. In 17% of companies, there are isolated cases, while another 17% suspect use but cannot prove it. The proportion of companies that categorically rule out shadow AI fell from 37% to 29%. The number of unreported cases is rising.

The DiWaBe 2.0 study (Germany, 2024) goes even further: 62% of employees use AI tools at work without official approval. Only about half of intensive AI users work with officially introduced applications.

And globally? An SAP study of 1,600 executives in eight markets shows that 64% of companies report that employees use unauthorized shadow AI tools at least occasionally. 77% expressed concerns about the associated risks.

Switzerland: High usage, little control

The situation in Switzerland deserves special attention. A KPMG study conducted in collaboration with the University of Melbourne (48,000 respondents in 47 countries) shows that 77% of Swiss respondents use AI in a professional context – well above the global average of 58%. That sounds progressive. The downside: over half of users admit to using AI contrary to company guidelines, for example by uploading sensitive data to freely accessible AI tools.

A Red Hat survey of 107 Swiss IT managers illustrates the extent of the problem: 97% say they are confronted with shadow AI. At the same time, only 3% of the Swiss companies surveyed are already generating measurable customer benefits from their AI investments.

The HWZ-Swisscom market study confirms that 38% of Swiss SMEs have introduced generative AI or are actively considering its use, but there is a lack of governance frameworks and clear guidelines almost everywhere.

Why do employees resort to shadow AI?

The motives are not malicious – they are pragmatic. The main drivers:

Productivity pressure: Employees are under time pressure and see AI tools as an immediate way to make their work easier. 68% of respondents to a Microsoft LinkedIn study said they struggle with the speed and scope of their work. AI becomes a "productivity prosthesis" that people no longer want to do without.

Lack of official offerings: Only 26% of German companies provide their own access to generative AI. The figure is even lower for SMEs. If no usable internal solutions are available, employees resort to private access. The consequence: uncontrolled proliferation.

Knowledge and governance gap: Around 73% of respondents feel poorly informed about AI. At the same time, the majority of organizations lack a clear AI policy. Less than half of Swiss employees have ever received specific AI training.

Fear of bans: Many employees use AI secretly because they fear sanctions or bureaucratic hurdles. In Switzerland, 69% of respondents say they pass off AI-generated content as their own work.

The risks: More than just an IT problem

Shadow AI is not purely a technical issue. The risks affect data protection, compliance, IT security, and ultimately the management personally.

Uncontrolled data leakage: Employees enter customer data, contracts, source code, or internal documents into external AI tools. This data ends up on third-party servers and can be stored, further processed, or used for training purposes. Companies lose control over the storage location, storage duration, and access.

GDPR violations: Shadow AI generally lacks a verified legal framework in accordance with Art. 6 GDPR – no valid consent, no data processing agreement, no documented balancing of interests. The accountability requirement under Art. 5(2) GDPR cannot be met because usage and data flows are neither documented nor auditable.

Third-country transfers: Many AI tools process data in cloud infrastructures outside the EU. Shadow AI circumvents standard contractual clauses and transfer impact assessments. This can lead to inadmissible third-country transfers – with corresponding fines of up to 4% of global annual turnover.

EU AI Act: With the gradual entry into force of the EU AI Act (2025–2027), the regulatory framework is becoming more stringent. Bans on certain AI practices have been in effect since February 2025. Requirements for AI competence and AI literacy are already in place. Anyone operating AI systems without sufficient control risks fines and liability consequences.

Cyber risks: CyberArk studies show that the growing prevalence of unauthorized AI applications is leading to an increase in unsecured machine identities with far-reaching access rights. 87% of respondents see identity silos as a direct cyber threat.

Swiss peculiarity: In the Swiss context, there are additional restrictive requirements for surveillance in the workplace (e.g., Art. 26 para. 1 ArGV 3), which further complicates the monitoring of shadow AI.

From bans to controlled integration

The studies are unanimous on this point: bans alone are ineffective. Employees circumvent bans if processes and tools do not follow suit. The recommended countermeasures are as follows:

Clear AI guidelines: Companies need a documented AI policy that regulates which tools may be used for which purposes, how data is classified, and which labeling requirements apply.

Secure internal platforms: Instead of five different AI services operating in the shadows, a shared platform is needed – ideally with European hosting and a clear data processing agreement.

Training and awareness: Employees need to understand how AI works, where the risks lie, and how to work with it safely. This applies to prompt hygiene, model limitations, and data protection requirements.

Establishing AI governance: Roles, approval processes, and risk assessments must be defined. The EU AI Act gradually requires a documented governance system for companies that use AI.

See shadow AI as an opportunity

And this is where it gets exciting for you as an entrepreneur or manager: shadow AI acts as a seismograph for productivity and innovation needs. Where shadow AI arises, there is a lack of transparent processes, secure alternatives, and comprehensible guidelines. Those who make secret AI use visible and integrate it in a controlled manner not only gain security, but also genuine process innovations.

The MIT study “State of AI in Business 2025” confirms that in over 90% of the companies surveyed, employees use private AI tools for work. The biggest investment therefore does not have to go into licenses, but into the organization's ability to learn. This is because the gap between companies that use AI systematically and those that remain in pilot project mode is widening.

And that's why I remain convinced that we can only do this together. Bans don't work, but with the right initiatives, you will also be able to counteract shadow AI.
www.rogerbasler.ch