revDSG & AI: The Shadow-AI Playbook for Swiss SMEs

Where uncontrolled AI use collides with the Swiss Data Protection Act — including personal liability — and the governance framework that fixes it.

The one-paragraph version: employees are already using AI tools — approved or not. Every paste of client data into a private ChatGPT or Claude account is data processing outside your contracts, records, and controls. Under the revised Swiss FADP (revDSG), that violates several articles at once, and unlike the GDPR, penalties of up to **CHF 250,000 can be levied against responsible individuals personally** (source: EDÖB 2024 report). The fix is not banning AI — it is making sanctioned AI easier and safer than shadow AI.

Where shadow AI breaks the revDSG

revDSG provision	How shadow AI violates it
Art. 7 — Privacy by design/default	Arbitrary data flows to unvetted platforms contradict the principle outright
Art. 9 — Processing contracts	Private AI accounts have no data-processing agreement with your company
Art. 16 — Cross-border transfers	Transfers to countries without adequacy need safeguards — shadow use bypasses them
Art. 12 — Records of processing	Processing your records don't know about = instant audit finding
Art. 22 — Impact assessments	High-risk processing needs a DPIA — impossible for tools you don't know exist

The personal dimension: FADP penalties target responsible individuals — executives, IT leads, compliance officers — up to CHF 250,000. For FINMA-regulated firms, shadow AI additionally collides with outsourcing and operational-risk requirements (Circular 2018/3).

The five zero-trust AI principles

- ☐ **Assume no inherent trust** — technical controls, not good intentions: block unapproved platforms, require authentication, log AI interactions
- ☐ **Verify continuously** — monitor what tools are used and what data they touch, not a one-time approval
- ☐ **Limit access explicitly** — least privilege for tools and for the data each tool can reach
- ☐ **Treat every AI agent as an identity** — inventoried, owned, governed like a human account
- ☐ **Policy-aware data access** — access checks before retrieval; sensitive data tagged and segmented

The 60–90 day implementation path

1. **Discovery (weeks 1–2):** network-level visibility of AI platform usage; anonymous employee survey (honesty beats blame); inventory of sanctioned tools; classify which data must never leave.
2. **Policy (weeks 3–4):** a short, written AI acceptable-use policy — approved tools, forbidden data categories, escalation path. One page beats twenty.
3. **Sanctioned alternative (weeks 4–8):** give teams a compliant AI setup that is genuinely better than the shadow tools — commercial API terms (no training on your data), Swiss/EU hosting options where residency is required, MCP connectors with read-only access instead of copy-paste.
4. **Controls & monitoring (weeks 8–12):** block-lists for unapproved platforms, logging for approved ones, quarterly review of usage patterns.

Your AI acceptable-use policy — the skeleton

- Approved tools and accounts (company accounts only — never private ones for work data)
- Data categories that must never reach external AI (client identities, health, financial records, credentials)
- Review duty: AI output is checked by a human before it leaves the company
- Who approves new tools, and how fast (a slow process re-creates shadow AI)
- What happens on violation — proportionate, known in advance

---

This guide is the condensed edition of our published in-depth article — sources and detail: **eflury.com/en/blog/shadow-ai-security-swiss-smes/**. As of July 2026. Not legal advice.

Questions about applying this in your business?

Free 30-minute call: eflury.com/en/contact/ · me@eflury.com · +41 79 910 77 87